

A remarkable ruby: Absorption in dense gas, rather than evolved stars, drives the extreme Balmer break of a little red dot at $z = 3.5$

Anna de Graaff^{1,*}, Hans-Walter Rix¹, Rohan P. Naidu^{2,**}, Ivo Labbé³, Bingjie Wang^{4,5,6}, Joel Leja^{4,5,6}, Jorryt Matthee⁷, Harley Katz⁸, Jenny E. Greene⁹, Raphael E. Hviding¹, Josephine Baggen¹⁰, Rachel Bezanson¹¹, Leindert A. Boogaard¹², Gabriel Brammer^{13,14}, Pratika Dayal¹⁵, Pieter van Dokkum¹⁰, Andy D. Goulding⁹, Michaela Hirschmann¹⁶, Michael V. Maseda¹⁷, Ian McConachie¹⁷, Tim B. Miller¹⁸, Erica Nelson¹⁹, Pascal A. Oesch^{20,13,14}, David J. Setton^{9,***}, Irene Shivaei²¹, Andrea Weibel²⁰, Katherine E. Whitaker^{22,13}, and Christina C. Williams²³

¹ Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

² MIT Kavli Institute for Astrophysics and Space Research, 77 Massachusetts Ave., Cambridge, MA 02139, USA

³ Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Melbourne, VIC 3122, Australia

⁴ Department of Astronomy & Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

⁵ Institute for Computational & Data Sciences, The Pennsylvania State University, University Park, PA 16802, USA

⁶ Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

⁷ Institute of Science and Technology Austria (ISTA), Am Campus 1, 3400 Klosterneuburg, Austria

⁸ Department of Astronomy & Astrophysics, University of Chicago, 5640 S Ellis Avenue, Chicago, IL 60637, USA

⁹ Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

¹⁰ Astronomy Department, Yale University, 219 Prospect St, New Haven, CT 06511, USA

¹¹ Department of Physics and Astronomy and PITT PACC, University of Pittsburgh, Pittsburgh, PA 15260, USA

¹² Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, The Netherlands

¹³ Cosmic Dawn Center (DAWN), Copenhagen, Denmark

¹⁴ Niels Bohr Institute, University of Copenhagen, Jagtvej 128, Copenhagen, Denmark

¹⁵ Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands

¹⁶ Institute of Physics, Laboratory for Galaxy Evolution, Ecole Polytechnique Fédérale de Lausanne, Observatoire de Sauverny, Chemin Pegasi 51, 1290 Versoix, Switzerland

¹⁷ Department of Astronomy, University of Wisconsin-Madison, 475 N. Charter St., Madison, WI 53706, USA

¹⁸ Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, 1800 Sherman Ave, Evanston, IL 60201, USA

¹⁹ Department for Astrophysical and Planetary Science, University of Colorado, Boulder, CO 80309, USA

²⁰ Department of Astronomy, University of Geneva, Chemin Pegasi 51, 1290 Versoix, Switzerland

²¹ Centro de Astrobiología (CAB), CSIC-INTA, Ctra. de Ajalvir km 4, Torrejón de Ardoz, E-28850 Madrid, Spain

²² Department of Astronomy, University of Massachusetts, Amherst, MA 01003, USA

²³ NSF's National Optical-Infrared Astronomy Research Laboratory, 950 North Cherry Avenue, Tucson, AZ 85719, USA

Received 21 March 2025 / Accepted 30 June 2025

ABSTRACT

The origin of the rest-optical emission of compact, red, high-redshift sources known as little red dots (LRDs) poses a major puzzle. If interpreted as starlight, it would imply that LRDs constitute the densest stellar systems in the Universe. However, alternative models suggest active galactic nuclei (AGN) may instead power the rest-optical continuum. Here, we present JWST/NIRSpec, NIRCам, and MIRI observations from the RUBIES and PRIMER programs of *The Cliff*: a bright LRD at $z = 3.55$ with an exceptional Balmer break, twice as strong as that of any high-redshift source previously observed. The spectra also reveal broad hydrogen ($H\alpha$ FWHM $\sim 1500 \text{ km s}^{-1}$) and He I emission, but no significant metal lines. We demonstrate that massive evolved stellar populations cannot explain the observed spectrum, even when considering unusually steep and strong dust attenuation or reasonable variations in the initial mass function. Moreover, the formally best-fit stellar mass and compact size ($M_* \sim 10^{10.5} M_\odot$, $r_e \sim 40 \text{ pc}$) would imply densities at which near-monthly stellar collisions might lead to significant X-ray emission. We argue that the Balmer break, emission lines, and $H\alpha$ absorption line are instead most plausibly explained by a black hole star (BH*) scenario, in which dense gas surrounds a powerful ionising source. In contrast to recently proposed BH* models of dust-reddened AGN, we show that spectral fits in the rest UV to near-infrared favour an intrinsically redder continuum over strong dust reddening. This may point to a super-Eddington accreting massive black hole or, possibly, the presence of (super)massive stars in a nuclear star cluster. *The Cliff* is the clearest evidence to date that at least some LRDs are not ultra-dense massive galaxies, and are instead powered by a central ionising source embedded in dense, absorbing gas.

Key words. galaxies: active – galaxies: evolution – galaxies: kinematics and dynamics – galaxies: stellar content

* Corresponding author: degraaff@mpia.de

** NASA Hubble Fellow.

*** Brinson Prize Fellow.

1. Introduction

Among the most debated discoveries accomplished with the James Webb Space Telescope (JWST) are the high-redshift compact red sources referred to as little red dots (LRDs). Although originally coined as a term for sources that were solely selected on their broad symmetric H α lines, they turned out to also be ubiquitously red and compact (Matthee et al. 2024), leading to a variety of definitions that are currently in use in the literature. Typically, photometric selections of LRDs require that sources are (1) compact in the rest-optical and (2) have v-shaped rest UV-to-optical spectral energy distributions (SEDs) or simply extremely red rest-frame optical colours (e.g. Labbé et al. 2025; Kocevski et al. 2025; Kokorev et al. 2024; Akins et al. 2024). Follow-up spectroscopy has revealed luminous and broad Balmer emission lines in a large fraction of these systems, as well as blue rest-UV and red rest-optical continua (Greene et al. 2024; Kocevski et al. 2025; Setton et al. 2024; Wang et al. 2024a; Hviding et al. 2025). The primary subject of debate now is whether these broad emission lines and v-shaped continua are powered by accreting massive black holes, likely constituting a new class of active galactic nuclei (AGN), or whether the SED properties can instead be attributed to massive stellar populations (e.g. Labbé et al. 2023; Greene et al. 2024; Wang et al. 2024a; Baggen et al. 2024; Guia et al. 2024; Mérida et al. 2025).

Determining whether the SEDs of LRDs are dominated by stars or AGN is a major challenge, even with high-quality rest-optical spectroscopy and photometry ranging from X-ray to radio wavelengths. Although broad (FWHM $\gtrsim 1000$ km s $^{-1}$) symmetric Balmer emission lines in combination with narrow forbidden lines are normally a key signature of broad-line AGN, LRDs lack other typical characteristics of AGN. Notably, LRDs show extremely weak, if any, X-ray emission (e.g. Ananna et al. 2024; Maiolino et al. 2024; Yue et al. 2024). Observations with JWST/MIRI have also revealed surprisingly blue colours in the rest-frame near- and mid-infrared (Wang et al. 2025; Williams et al. 2024; Pérez-González et al. 2024; Setton et al. 2025), which implies a lack of hot dust emission from AGN tori. Moreover, while some LRDs do show AGN signatures at rest-UV wavelengths in the form of strong emission lines (Labbé et al. 2024; Akins et al. 2025a), the UV emission of many others can be well modelled by a star-forming galaxy (e.g. Maiolino et al. 2024; Killi et al. 2024; Barro et al. 2024a).

The origin of the rest-optical continuum emission is perhaps the most contentious. If dominated by the emission from a strongly dust-attenuated AGN, the implied host galaxy mass is low, resulting in a black hole-to-host mass ratio that lies far above the well-established scaling relations in the local Universe (e.g. Maiolino et al. 2024; Pacucci et al. 2023; Furtak et al. 2024), although this may in part be driven by uncertainties in the black hole and galaxy masses as well as selection effects (e.g. Li et al. 2025; Lupi et al. 2024; Wang et al. 2024a). If confirmed, such a large population of overmassive black holes may have important implications for the seed masses and growth mechanisms of early black holes (e.g. Kokorev et al. 2023; Bogdán et al. 2024; Goulding et al. 2023), and would be in tension with current cosmological hydrodynamical simulations (e.g. Habouzit 2025; Matthee et al. 2025). Intriguingly, however, the transition in the v-shaped SEDs from blue to red has been shown to occur at the Balmer limit for $\sim 50\%$ of LRDs (Setton et al. 2024), and some LRDs show Balmer breaks that are as strong as those observed in high-redshift post-starburst galaxies (e.g. Labbé et al. 2024; Wang et al. 2024a; Ma et al.

2025). At face value, it therefore appears reasonable to interpret the rest-optical as being dominated by emission from an evolved stellar population in a host galaxy instead of by an AGN.

However, such a stellar mass-dominated interpretation leads to great tension with galaxy formation models. The stellar masses of LRDs obtained from SED fitting are typically high ($M_* \sim 10^{10} M_\odot$ under the assumption of a Milky Way-like initial mass function; e.g. Akins et al. 2024), with some galaxies being as massive as $M_* \sim 10^{11} M_\odot$ by $z \sim 7-8$ (Labbé et al. 2023; Wang et al. 2024a). Coupled with the relatively high number densities of LRDs (e.g. Kokorev et al. 2024, and thus modest implied halo masses, e.g. Pizzati et al. 2025), this would imply very high conversion rates from available baryons in haloes into stellar mass. The most luminous LRDs are estimated to be so massive that the extreme galaxy formation efficiencies required for their formation are difficult to accommodate within the Λ CDM model (Boylan-Kolchin 2023). Moreover, because of the very compact morphologies of LRDs, the high masses also imply exceptionally high stellar mass densities that exceed the maximum expected stellar densities from observations of dense star clusters and star formation models (Baggen et al. 2023, 2024; Guia et al. 2024; Labbé et al. 2024; Ma et al. 2025).

To bridge the predominantly AGN or predominantly stellar scenarios, SED modelling efforts have therefore focused on multi-component fitting of photometric (e.g. Barro et al. 2024a; Pérez-González et al. 2024; Juodžbalis et al. 2024; Leung et al. 2024) and spectrophotometric data (Harikane et al. 2023a; Labbé et al. 2024; Ma et al. 2025; Maiolino et al. 2024; Wang et al. 2025, 2024a). Generally, including a dust-reddened AGN component lowers the inferred stellar mass and stellar mass density (by up to ~ 2 dex), reducing some of the aforementioned tension with galaxy formation models. However, several of these SED models, especially when fitting spectra, require an uncomfortable level of tuning; for example, strong AGN emission must consistently arise just redward of the Balmer break, and an unusually steep dust extinction law is needed to match the spectral shape (Wang et al. 2024a; Ma et al. 2025).

Most recently, the interpretation that the observed Balmer break must originate from stellar emission has been called into question (Inayoshi & Maiolino 2025; Ji et al. 2025). Instead, these studies suggest that extremely dense gas in the close vicinity of an AGN can resemble the conditions of ($\sim 10^4$ K) stellar atmospheres, and hence result in spectra that mimic the SEDs of high-redshift quiescent galaxies (although the AGN models of Ji et al. 2025 require substantially greater dust attenuation). This is further motivated by the discovery of strong absorption features in the Balmer and He I emission lines, which point to the presence of absorbing dense gas clouds (Matthee et al. 2024; Wang et al. 2025; Juodžbalis et al. 2024; D'Eugenio et al. 2025). Moreover, such an AGN-dominated scenario would also be in line with recent clustering measurements that independently suggest low host galaxy masses in a sample of fainter LRDs (Matthee et al. 2025).

Critically, however, to date it has been difficult to robustly differentiate between these various models, especially for the brightest LRDs with strong Balmer breaks. That is, models in which a massive galaxy dominates the SED often fit the rest-frame UV to near-IR spectrophotometric data equally well as multi-component galaxy + AGN models or the newly proposed non-stellar Balmer break model. Arguments against or in favour of specific models have therefore relied on other evidence, such as the X-ray non-detections, emission line kinematics, morphology, or the degree of model fine-tuning required.

In particular, conclusively ruling out the presence of a massive, ultra-dense stellar population dominating the rest-optical has been surprisingly difficult. Although this scenario requires invoking high star formation efficiencies, its simplicity is appealing. The SED, including the lack of X-ray emission and hot dust in the mid-IR, can largely be explained with standard stellar population models, and the broad emission line kinematics have been suggested to arise naturally from the very high stellar mass density (Baggen et al. 2024).

In this paper we present a luminous LRD at $z_{\text{spec}} = 3.548$ from the Red Unknowns: Bright Infrared Extragalactic Survey (RUBIES; de Graaff et al. 2025a), named *The Cliff*, as JWST/NIRSpec spectroscopy reveals an exceptionally strong Balmer break. We rigorously assessed the possibility that this remarkable source is stellar in origin by testing a range of stellar and AGN models and by exploring its implied stellar dynamical properties. We demonstrate that we can, for the first time, robustly exclude a high stellar mass and high stellar density origin for a LRD with a strong Balmer break. The photometric and spectroscopic data are described in Section 2, and the measurements of the emission line and the morphological properties are presented in Sections 3 and 4, respectively. We present our stellar population modelling in Section 5, and explore the consequences of these models on the dust properties and galaxy dynamical properties in Sections 6.1 and 6.2. We also assess possible variations in the stellar initial mass function in Section 6.3. Finally, we propose in Section 6.4 that the preferred model for this source is likely that of an accreting massive black hole embedded in high-density gas, and then discuss the intrinsic spectrum of the ionising source. Our observations and findings are summarised in Section 7. Where relevant, we adopt the best-fit cosmological parameters from the WMAP 9 yr results: $H_0 = 69.32 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.2865$, and $\Omega_\Lambda = 0.7135$ (Hinshaw et al. 2013).

2. Data

2.1. Imaging

The primary object of this paper lies in the UDS field (RA 02:17:38.58; Dec −05:07:46.79), and was selected (see Section 2.2) from public JWST/NIRCam and MIRI imaging obtained as part of the PRIMER Survey (GO-1837, PI: Dunlop). The PRIMER imaging covers 8 NIRCam filters (F090W, F115W, F150W, F200W, F277W, F356W, F410M, F444W; see e.g. Donnan et al. 2024 for details) as well as 2 MIRI filters (F770W, F1800W).

We use the publicly available image mosaics from the Dawn JWST Archive (DJA, version 7.2), which were reduced with the *grizli* software (Brammer 2023) as described in Valentino et al. (2023). These mosaics have a pixel scale of $0.04''$; for the purposes of morphological analysis in Section 4, we also use a custom image mosaic for the F200W filter with a better pixel sampling of $0.02''$ (as was also used in Weibel et al. 2025).

The inset of Figure 1 shows a colour image from the F115W, F277W, and F444W filters, revealing a compact red source (our target) as well as a blue neighbouring source. We obtain a photometric redshift $z_{\text{phot}} = 2.78^{+0.17}_{-0.05}$ of this neighbour from the catalogue of Weibel et al. (2024), substantially lower than the redshift of *The Cliff*. Although it is not physically associated, we do find that the circular aperture photometry at short wavelengths is contaminated by the foreground source, because the foreground source has a very blue colour and the photometry was measured

from mosaics that were matched to the coarser spatial resolution of the F444W filter.

We therefore perform simultaneous Sérsic profile fitting with *pysersic* (Pasha & Miller 2023) in order to obtain deblended fluxes. Each source is modelled as a single Sérsic profile, and convolved with the empirical point spread functions (PSFs) of Weibel et al. (2024). We fit the available NIRCam filters using uniform priors for all parameters, restricting the Sérsic index to the range $[0.65, 4]$ and the effective radius (major axis) to $[0.5, 100] \text{ pix}$, and using the No U-turn sampler of *numpyro* to sample the posterior distributions (Hoffman & Gelman 2014; Phan et al. 2019). The NIRCam images, maximum a posteriori (MAP) models, and residual images can be found in Appendix A. As described in Appendix A, we find that the Sérsic profile model of *The Cliff* strongly dominates the flux at the position of the NIRSpec microshutter. In other words, contamination from the nearby neighbour does not significantly contribute to the observed NIRSpec spectrum or affect any of the conclusions in this work.

The MIRI images are shallower and we therefore only fit for the flux, informed by the fit in the NIRCam/F200W filter and using the empirical MIRI PSF models of Libralato et al. (2024). We model the foreground source as a single Sérsic profile, and set truncated (at $\pm 3\sigma$) Gaussian priors for the model parameters based on the posteriors of the F200W fit. Because the posterior of the effective radius of *The Cliff* is at the border of the prior ($r_{\text{e,maj}} \approx 0.5 \text{ pix}$), we fit a simpler point source model instead of a Sérsic model for the MIRI filters (with a prior on the position from the F200W posteriors). The pixel scale of the MIRI image mosaics ($0.04''$) is substantially smaller than the pixel scale of the MIRI instrument ($0.11''$), and the pixels are therefore highly correlated. This correlated noise is not accounted for by *pysersic*, which only uses the variance image. To estimate uncertainties on the derived fluxes, we therefore perform point source photometry with *pysersic* in random empty areas in the vicinity of *The Cliff*, masking bright sources. We take the standard deviation of these random measurements as the uncertainty on the point source flux, which is approximately two times higher than the formal uncertainties of *pysersic*. A table with all NIRCam and MIRI fluxes, as well as the images, morphological models and residual images, are provided in Appendix A.

2.2. Spectroscopy

Observations with the JWST/NIRSpec microshutter array (MSA) were obtained in July 2024 as part of RUBIES (GO-4233; PIs: de Graaff & Brammer). *The Cliff* (RUBIES-UDS-154183) was selected with high priority for its very red colour (F150W–F444W ≈ 3.9) and promoted to the highest Priority Class 0 because of its extremely compact morphology and strong photometric Balmer break.

The details of the RUBIES priority classes and observing strategy can be found in de Graaff et al. (2025a). Briefly, we observe with both the low-resolution PRISM/Clear ($0.6\text{--}5.3 \mu\text{m}$) and medium-resolution G395M/F290LP ($2.9\text{--}5.2 \mu\text{m}$) disperser/filter combinations, using a 3-point nodding pattern and total exposure time of 48 min per disperser. The data were reduced with the *msaexp* software and correspond to version 3 of the DJA, as described in Heintz et al. (2025) and de Graaff et al. (2025a). Notably, this version provides two types of background subtraction for the PRISM spectra: a local subtraction from the image differences between the nodded exposures, and a global background subtraction obtained from empty sky shutters in the mask. For consistency, because only the local background

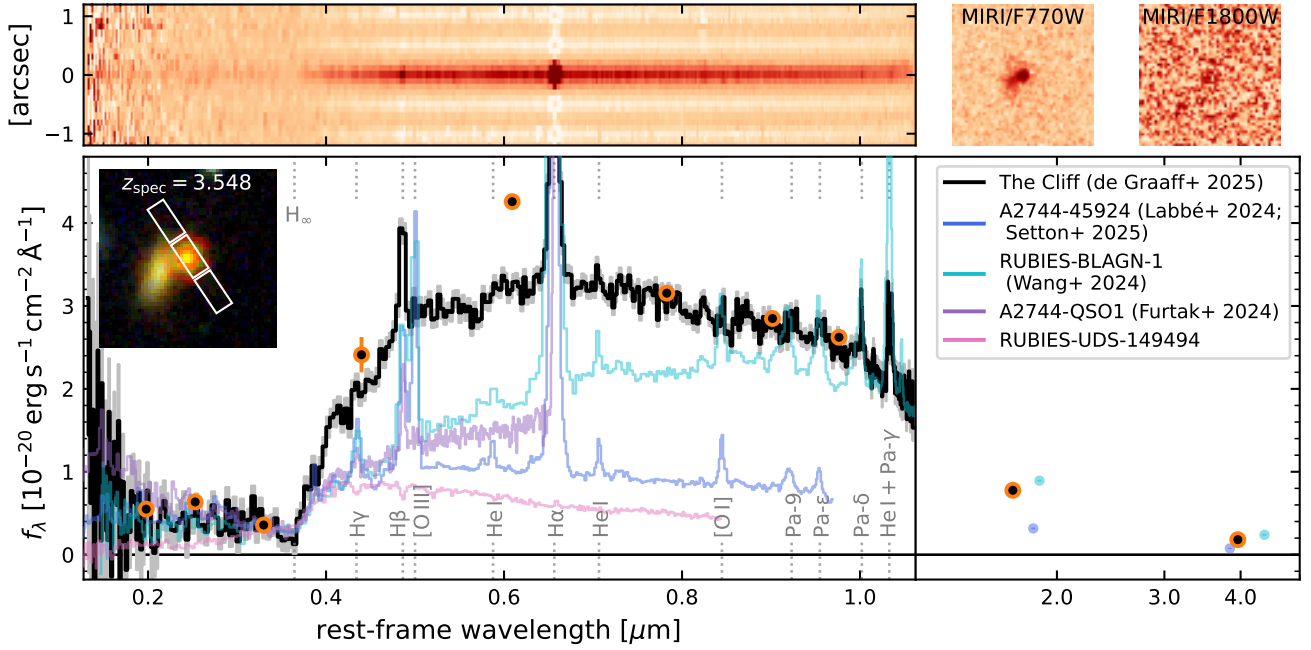


Fig. 1. NIRSpect/PRISM spectrum of *The Cliff* (RUBIES-UDS-154183) at $z_{\text{spec}} = 3.548$. The orange points show NIRC and MIRI photometry spanning an observed wavelength range of 0.9–18 μm ; the inset colour image was constructed using the NIRC F115W, F277W and F444W filters and shows the location of the NIRSpect microshutters. The coloured lines are the NIRSpect/PRISM spectra (rescaled using the median flux at rest-frame [3200–3700] \AA) of four sources with strong Balmer breaks: two of the most luminous LRDs in the literature ($z_{\text{spec}} \sim 3\text{--}5$; Labbe et al. 2024; Wang et al. 2025) and their MIRI detections (Setton et al. 2025), the triple-imaged LRD at $z_{\text{spec}} = 7.04$ of Furtak et al. (2024), and a massive post-starburst galaxy at $z_{\text{spec}} = 4.62$ (a medium-resolution NIRSpect spectrum of which was presented by Carnall et al. 2024). The dotted lines indicate the locations of strong emission line features in the different spectra as well as the Balmer limit. *The Cliff* shows an exceptionally strong Balmer break, a declining SED in the rest near-infrared ($\sim 1\text{--}4 \mu\text{m}$), strong H and He I emission lines, but (in comparison to the two luminous LRDs) very weak metal lines.

subtraction is available for the G395M spectra, we use the local background subtraction for both dispersers throughout (although we note that in practice this makes no significant difference to any of the results in this paper).

2.3. X-ray

We obtain public Chandra X-ray data from the X-UDS survey (Kocevski et al. 2018). The location of *The Cliff* was observed with effective exposure times of 573.9 ks ($< 2 \text{ keV}$) and 609.2 ks ($2\text{--}7 \text{ keV}$). Following Wang et al. (2025), we obtain 3σ upper limits on the X-ray luminosities: $L_X < 1.9 \times 10^{44} \text{ erg s}^{-1}$ (rest $1\text{--}4 \text{ keV}$), $L_X < 3.5 \times 10^{42} \text{ erg s}^{-1}$ (rest $4.5\text{--}9 \text{ keV}$), $L_X < 6.7 \times 10^{43} \text{ erg s}^{-1}$ (rest $10\text{--}30 \text{ keV}$).

3. Spectral properties

The 2D PRISM spectrum of *The Cliff* and its 1D extraction are shown in Figure 1. The NIRC and MIRI fluxes are also overplotted, and demonstrate that the absolute flux calibration of the NIRSpect spectrum is in good agreement with the photometry. We further show the rescaled NIRSpect/PRISM spectra of four sources at $z_{\text{spec}} > 3$ with very strong Balmer breaks: A2744-45924 ($z_{\text{spec}} = 4.47$; Labbe et al. 2024) and RUBIES-BLAGN-1 ($z_{\text{spec}} = 3.10$; Wang et al. 2025) are two of the most luminous LRDs found to date; the triply imaged LRD A2744-QSO1 has the strongest Balmer break observed so far at $z > 6$ ($z_{\text{spec}} = 7.04$; Furtak et al. 2024); RUBIES-UDS-149494 is a massive quiescent galaxy at $z_{\text{spec}} = 4.62$ (a medium-resolution spectrum of which was also presented in Carnall et al. 2024

with ID PRIMER-EXCELS-117560). Comparing to the other LRDs, *The Cliff* also shows strong Balmer and Paschen emission lines, and a blue SED shape in the rest-frame near-infrared (Setton et al. 2025). Similar to A2744-QSO1, but unlike the most luminous LRDs, *The Cliff* does not have strong [O III] $\lambda 5008$ emission or any other strong metal lines (further quantified in Section 3.2).

3.1. Balmer break

Most strikingly, *The Cliff* stands out for its extremely strong Balmer break in comparison to the four literature sources. We quantify this in Figure 2, where we compile a large sample of LRDs (Furtak et al. 2024; Labbe et al. 2024; Setton et al. 2024; Wang et al. 2024a) and massive quiescent galaxies at $z > 3$ (Barrufet et al. 2025; Carnall et al. 2024; de Graaff et al. 2025b; Glazebrook et al. 2024; Nanayakkara et al. 2025; Weibel et al. 2025) with public JWST/NIRSpect spectroscopy obtained from the DJA. To measure the Balmer break strength we integrate the spectrum in two tophat filters with wavelength ranges [3620, 3720] \AA and [4000, 4100] \AA using pyphot (Fouesneau 2025), and compute the flux density ratio (in f_v)¹. Error bars are estimated using 500 random Gaussian draws from the error spectrum. For reference, we also show the Balmer break

¹ This definition differs slightly from that by Wang et al. (2025), who used the same wavelength ranges, but computed the ratio of the median flux density in f_λ . However, the definition used here more closely matches the commonly measured Dn4000 index of Balogh et al. (1999), and also allows for better estimation of the uncertainty.




Fig. 2. Balmer break strength, measured as the ratio of the mean flux density in the rest-frame wavelength ranges $[3620, 3720] \text{ \AA}$ and $[4000, 4100] \text{ \AA}$ from public JWST/NIRSpec data. Grey squares show the median values of stacks of star-forming galaxies from Roberts-Borsani et al. (2024); red markers represent a compilation of LRDs (Furtak et al. 2024; Labbe et al. 2024; Setton et al. 2024; Wang et al. 2024a); blue markers show a large sample of massive quiescent galaxies (compiled from Barrufet et al. 2025; Carnall et al. 2024; de Graaff et al. 2025b; Glazebrook et al. 2024; Nanayakkara et al. 2025; Weibel et al. 2025). The recently discovered LRD of Naidu et al. (2025), MoM-BH*-1, is a higher-redshift analogue of *The Cliff* and discussed in Section 6.4.

strengths measured from stacks of star-forming galaxies by Roberts-Borsani et al. (2024).

The Balmer break of *The Cliff* ($6.9^{+2.8}_{-1.5}$) is ≥ 2 times stronger than that of any high-redshift massive quiescent galaxy (< 3.1), and all high-redshift LRDs with Balmer breaks published to date (the largest break strength being ≈ 3.56 for A2744-45924 of Labbe et al. 2024). The only source (MoM-BH*-1; Naidu et al. 2025) that matches the break strength of *The Cliff* is a LRD that was discovered at approximately the same time in JWST Cycle 3 program GO-5224 (PIs: Oesch & Naidu), and is discussed in detail in Section 6.4.

3.2. Emission lines

The spectra reveal a suite of emission lines, primarily Balmer, Paschen and He I lines. The G395M spectrum covers the rest-frame wavelength range of $0.64\text{--}1.14 \text{ \mu m}$ ($2.9\text{--}5.2 \text{ \mu m}$ observed; with the chip gap falling at $4.6\text{--}4.8 \text{ \mu m}$), and therefore kinematically resolves the $H\alpha$ line as well as various Pa features. Figure 3 shows zoomed in wavelength ranges of the strongest emission lines of both the PRISM and G395M spectra.

The $H\alpha$ emission shows a complex, broad and asymmetric profile. We attribute this asymmetry to a redshifted Balmer absorption line, a feature that has been observed in several other LRDs and likely originates from absorbing dense gas clouds along the line of sight (e.g. Matthee et al. 2024; Judžbalis et al.

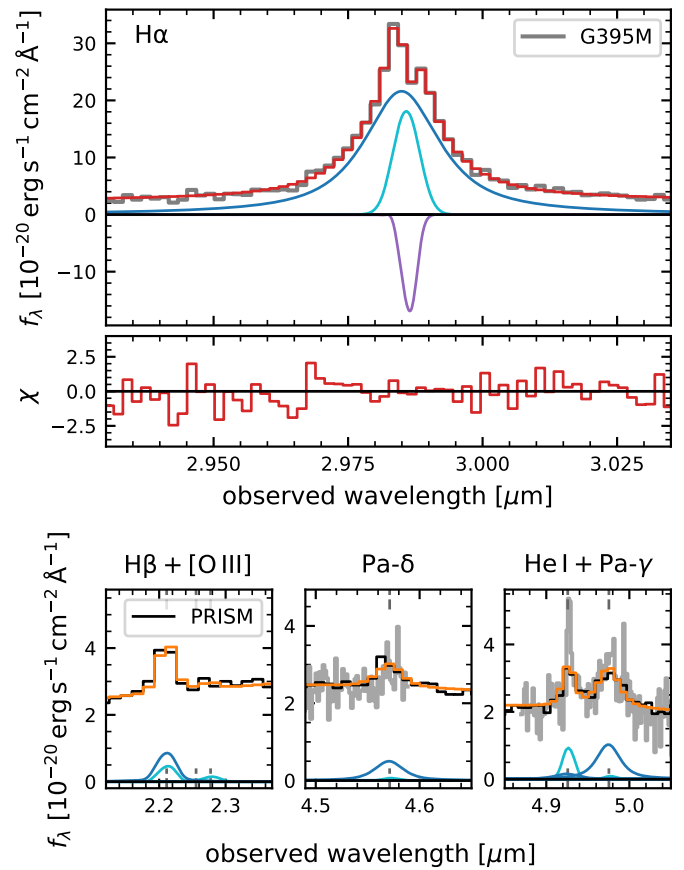


Fig. 3. Kinematic emission line decomposition from simultaneous fitting to PRISM (black) and G395M (grey) spectra. Top: $H\alpha$ emission and absorption complex. It is well-described by a broad Lorentzian profile (blue; $\text{FWHM} \sim 1400 \text{ km s}^{-1}$), with weaker narrow emission (cyan) as well as redshifted absorption (purple). The total G395M model and residuals are shown in red. Bottom: Zoomed-in images of other strong H and He features, also revealing a non-detection of the $[O III]$ doublet (see Table 1). Although (where possible) both dispersers were used in the fitting, in the bottom panels we show only the PRISM model components and total model (orange).

2024; D'Eugenio et al. 2025). It is likely that this absorption is also present in other lines (e.g. He I or $H\beta$; Wang et al. 2025; Ji et al. 2025), but at the current signal-to-noise of the data (S/N) and available spectral resolution we cannot establish whether this is the case.

We therefore begin by focusing on the high S/N $H\alpha$ line complex in the G395M spectrum, using the Bayesian emission line fitting software unite described in Hviding et al. (2025). Briefly, this builds on our prior fitting methods (as in Wang et al. 2025, 2024a; de Graaff et al. 2025a) and therefore robustly accounts for the undersampling of the JWST/NIRSpec line spread function (LSF) through integration of the emission line profiles. The error spectra are automatically rescaled to match the measured scatter in the continuum blue- and red-ward of the emission line. Moreover, our fitting accounts for calibration uncertainty in the NIRSpec LSF: we assume the LSF of an idealised point source from de Graaff et al. (2024), but introduce a nuisance parameter f_{LSF} by which the dispersion is multiplied to model both the systematic and measurement uncertainty in the LSF; the prior for this nuisance parameter is a truncated Gaussian centred at $f_{\text{LSF}} = 1.2$ and dispersion of 0.1, with minimum and maximum values of $[0.9, 1.5]$. As in Section 2.1, the

probabilistic model is implemented in `numpyro` using the No U-Turn Sampler to generate posterior samples.

Our fiducial model for the $H\alpha$ complex, presented in Figure 3 and Table 1, assumes a broad Lorentzian and narrow Gaussian emission component, a narrow Gaussian absorption component, and a linear continuum. The redshifts of all components are allowed to vary; we use uniform priors for the line widths and fluxes. We have tested a range of model variations, which we do not show but describe here, performing model comparison to the fiducial model using the Widely Applicable Information Criterion (WAIC; Watanabe & Oppen 2010, although the Bayesian Information Criterion yields identical conclusions). First, we find that a broad Lorentzian component provides a significantly better fit over a broad Gaussian component² ($\Delta\text{WAIC} = 39.8$), as was also found for the broad $H\alpha$ line of A2744-45924 (Labbe et al. 2024). Second, although the flux of the absorption component is poorly constrained in our fiducial model, the model with an absorption component is formally strongly preferred over one without ($\Delta\text{WAIC} = 27.6$). Deep, high-resolution observations would be needed to robustly constrain the kinematics and equivalent width of the absorber. Finally, we have tried to include the [N II] doublet in the fit (with fixed flux ratio of 1:2.95 for the doublet). This fit is weakly disfavoured compared to the fiducial model ($\Delta\text{WAIC} = 3.76$), and the flux of the [N II] doublet is consistent with zero and all other parameters are consistent with the values presented in Table 1.

Next, we perform simultaneous fitting to the PRISM and G395M spectra to obtain emission line properties for a broader suite of lines. The key benefit of this approach is that it leverages the S/N of both spectra. We introduce two new parameters in order to enable simultaneous fitting of multiple dispersers, which are necessary due to systematic uncertainties in the flux and wavelength calibration. Specifically, we fit for a flux offset between the G395M and PRISM dispersers as well as a detector pixel offset, and set Gaussian priors on these parameters based on the population average offsets from a large number of spectra (for details on calibration issues and typical offsets, see de Graaff et al. 2025a). We use the following setup for the model: we assume a Lorentzian broad component, as well as a Gaussian narrow emission component for the $H\beta$, $H\alpha$, Pa- δ , Pa- γ , He I $\lambda 10830$, and O I $\lambda 8446$ lines. For $H\alpha$ we also add a Gaussian absorption component. We further model the [O III] $\lambda\lambda 4960, 5008$ emission lines with only a narrow Gaussian component, fixing the flux ratio of the doublet to 1:2.98. Due to the modest S/N for most emission lines, all narrow and broad line velocity widths are tied together; the broad and narrow line redshifts are allowed to deviate. We further assume a linear continuum in the region around each emission line ($\pm 15\,000\text{ km s}^{-1}$).

The line fluxes from this joint fit are presented in Table 1, and best-fit models are shown for the G395M (red) and PRISM (orange) spectra in Figure 3. The narrow and broad component redshifts and velocity dispersions of this joint fit are consistent with our previous fit to the $H\alpha$ line alone, and we therefore do not list these separately. The [O III] and O I emission lines are not significantly detected, which may imply a low metallicity or high gas density (see Section 6.4), and we therefore provide the 95th percentiles as upper limits. The total $H\beta$ emission is detected at the $\sim 10\sigma$ level, but its broad and narrow components are indi-

Table 1. Emission line properties.

Line fluxes		
Line	Flux (narrow) ($10^{-18}\text{ erg s}^{-1}\text{ cm}^{-2}$)	Flux (broad) ($10^{-18}\text{ erg s}^{-1}\text{ cm}^{-2}$)
$H\beta$	$1.7^{+1.7}_{-1.2}$	$3.8^{+1.6}_{-2.0}$
[O III] $\lambda\lambda 4960, 5008$	<1.5	
$H\alpha$ (emission)	$13.0^{+4.8}_{-3.0}$	$61.6^{+1.8}_{-2.0}$
$H\alpha$ (absorption)	$-7.2^{+2.2}_{-5.1}$	
O I $\lambda 8446$	<0.5	<0.9
Pa- δ	$0.13^{+0.16}_{-0.09}$	$2.4^{+0.4}_{-0.4}$
Pa- γ	$0.15^{+0.18}_{-0.11}$	$5.2^{+0.4}_{-0.5}$
He I $\lambda 10830$	$1.8^{+0.3}_{-0.3}$	$0.9^{+0.5}_{-0.5}$
$H\alpha$ kinematic properties		
z_{narrow}	$3.5484^{+0.0004}_{-0.0004}$	
z_{broad}	$3.5470^{+0.0003}_{-0.0003}$	
Δv_{absorb}	$49^{+26}_{-21}\text{ km s}^{-1}$	
$\text{FWHM}_{\text{narrow}}$	$478^{+95}_{-93}\text{ km s}^{-1}$	
$\text{FWHM}_{\text{broad}}$	$1533^{+110}_{-80}\text{ km s}^{-1}$	
$\text{FWHM}_{\text{absorb}}$	$133^{+95}_{-92}\text{ km s}^{-1}$	

vidually poorly constrained due to the low spectral resolution of the PRISM at $\sim 2\mu\text{m}$ ($R \sim 50$).

4. Morphology

We previously (Section 2.1) established that *The Cliff* has a compact morphology. However, it is challenging to quantify robustly whether a LRD is extremely compact or truly unresolved, as many high-redshift LRDs are very faint in the short wavelength NIRCcam filters that have the highest spatial resolution. Only for a few objects has it been possible to obtain robust upper limits on the size based on such rest-UV imaging (e.g. Furtak et al. 2024, upper limit of $<30\text{ pc}$). Others have typically either assumed conservative upper limits on the size (e.g. the half width at half maximum at $4\mu\text{m}$ of $\lesssim 400\text{ pc}$, Leung et al. 2024), or focused on Sérsic profile fitting in the long wavelength NIRCcam filters (e.g. Baggen et al. 2023, 2024; Akins et al. 2024, resulting in size upper limits of $\lesssim 100\text{--}200\text{ pc}$). Because of its lower redshift, the target of this paper is fortunately bright also in the reddest short-wavelength NIRCcam filter (F200W, probing rest-frame $0.44\mu\text{m}$) for which the PSF FWHM $\approx 0.066''$.

The second uncertainty in modelling compact morphologies is the PSF model itself, which is typically constructed empirically from a stack of stars selected across the entire mosaic. Therefore, it represents an average PSF that may deviate in detail from the true local PSF due to small geometric distortions. Here we use the empirical PSFs constructed using the methods described in Weibel et al. (2024), with stars selected in the magnitude range F200W $\sim 19\text{--}24.5$.

Serendipitously, a star of F200W = 25.1 lies $\sim 1''$ away from *The Cliff* (see Figure 4). This star is too faint to have been included in our empirical PSF model, but it is approximately equally bright as the primary target (for which F200W = 25.3). In other words, the nearby star provides an ideal test of our empirical PSF model and Sérsic profile fitting.

We perform simultaneous fitting with `pysersic` as already described in Section 2.1, modelling the foreground source, star, and LRD with single Sérsic profiles, but now for the F200W

² This difference in modelling likely also largely explains the difference between the broad $H\alpha$ FWHM obtained here and the larger FWHM measured by Taylor et al. (2025) for *The Cliff* (FWHM $\sim 2300\text{ km s}^{-1}$), who used a two-component Gaussian model.

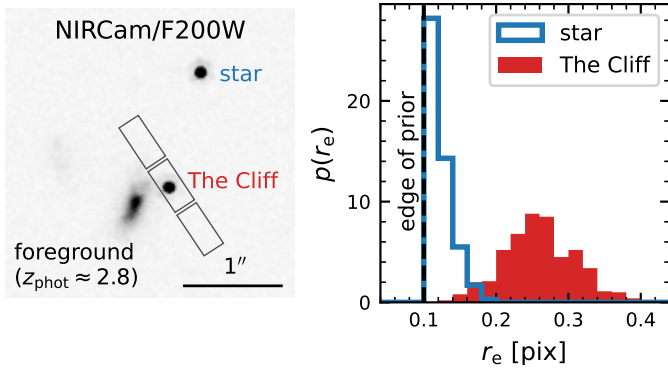


Fig. 4. Left: NIRCcam/F200W image ($0.02'' \text{ pix}^{-1}$) of *The Cliff*, the nearby foreground source at a lower redshift of $z_{\text{phot}} \approx 2.8$, and a star of approximately equal magnitude (at a distance of $\sim 1''$). All sources are modelled as single Sérsic profiles and fit simultaneously to account for blending. Right: Posterior distribution of the effective radius of the star is clustered at the edge of the prior, as expected for an unresolved point source, while that of the LRD converges to a small value of $r_{e,\text{maj}} \approx 0.26 \text{ pix} \approx 40 \text{ pc}$.

mosaic with a pixel scale of $0.02''$ (see Appendix A for the MAP model and residual image). This time, the prior on the Sérsic index $([0.65, 6.0])$ and effective radius $([0.1, 100] \text{ pix})$ are allowed to be broader. We find that if we use a lower bound of $r_{e,\text{maj}} = 0.5 \text{ pix}$, the posterior distributions of the star and LRD are identical for the size, axis ratio and Sérsic index. However, when we lower this bound to 0.1 pix these posteriors begin to deviate. The axis ratios and Sérsic indices are consistent within the uncertainties (Table 2), but the sizes differ significantly. In Figure 4 we show that the posterior of the effective radius piles up at the boundary of the prior for the star, which is expected for an unresolved point source, and therefore can be interpreted as a success of our empirical PSF model. On the other hand, the fit to the LRD converges to a very small value of $r_{e,\text{maj}} \sim 0.26 \text{ pix}$, corresponding to a physical major axis size of $\sim 40 \text{ pc}$.

The LRD therefore could formally be considered to be marginally resolved. However, it is also possible that a bright point source dominates over a faint (more extended) host galaxy, leading to a slightly resolved size when fitting a single Sérsic component. We further caution that the nearby foreground source appears to have diffuse extended emission, and its morphology deviates from a simple, symmetric Sérsic profile. This residual flux may therefore bias the inferred size of the LRD to higher values. Nevertheless, our Sérsic profile fitting does provide a stringent upper limit on the size of $<51 \text{ pc}$ (95th percentile; or $<58 \text{ pc}$, 99.7th percentile), which is substantially lower than most upper limits for unlensed LRDs in current literature.

5. Stellar population modelling

We now turn to the physical interpretation of the SED, and ask whether the observed Balmer break and spectral shape could be explained by conventional stellar population models. A range of models and codes have been used in recent literature to fit the SEDs of LRDs, with varying degrees of success (as described in Section 1). Here, we focus only on tools that are able to fit spectroscopic data and on models that are able to produce a sharp Balmer break. Specifically, we apply the modelling of Wang et al. (2025, 2024a) and Labbe et al. (2024), but do not consider their AGN-only or AGN-dominant models, as those

Table 2. Morphological properties.

Source	$r_{e,\text{maj}}$ (pix)	$r_{e,\text{maj}}$ (pc)	axis ratio	Sérsic index
star	$0.117^{+0.021}_{-0.011}$	$17.5^{+3.2}_{-1.7}$	$0.41^{+0.30}_{-0.22}$	$0.92^{+0.32}_{-0.19}$
LRD	$0.259^{+0.050}_{-0.046}$	$38.6^{+7.4}_{-6.9}$	$0.47^{+0.29}_{-0.26}$	$1.17^{+0.42}_{-0.30}$

were already shown to fail to produce strong enough Balmer breaks for LRDs that are less extreme than *The Cliff*.

5.1. Prospector models

Following the methodology described in Wang et al. (2025, 2024a), we jointly fit the PRISM spectrum and NIRCcam+MIRI photometry with the Bayesian modelling software Prospector (Leja et al. 2017; Johnson et al. 2021), using the dynamic nested sampling code dynesty (Speagle 2020). Model spectra are redshifted using the best-fit redshift of Section 3.2 and convolved to the NIRSpec/PRISM resolution using the model LSF of an idealised point source of de Graaff et al. (2024). To account for small mismatches between the flux calibration of the spectrum and photometry, we allow for a constant rescaling (i.e. a zeroth-order polynomial calibration vector) between the spectrum and the photometry. We find that higher-order polynomials ($n \geq 2$) lead to calibration vectors that are significantly model-dependent in shape and amplitude, indicating that this vector corrects for the mismatch between model and data rather than any real flux offset between the spectrum and photometry. Emission lines are not interpreted physically in the modelling, i.e. we fit broad Gaussians for the $H\beta$, $H\alpha$, $\text{Pa-}\delta$, $\text{Pa-}\gamma$, and He I lines with flux ratios that are allowed to vary arbitrarily.

5.1.1. Galaxy-only model

We begin by fitting a model that includes a stellar population component as well as emission lines, but no AGN component: this is the galaxy-only model of Wang et al. (2025), and uses the model setup and informative priors of the Prospector- β model described in Wang et al. (2024b). Briefly, this uses the Flexible Stellar Population Synthesis models (FSPS; Conroy & Gunn 2010), with MIST stellar isochrones (Choi et al. 2016; Dotter 2016) and the MILES stellar library (Sánchez-Blázquez et al. 2006), and assumes a Chabrier (2003) initial mass function (IMF; the possibility of IMF variations is discussed in Section 6.3). The star formation history follows a modified version of the non-parametric Prospector- α model of Leja et al. (2017) with seven logarithmically spaced time bins that are optimised for the high-redshift Universe. Dust attenuation follows a two-component dust screen model (Charlot & Fall 2000) with the flexible attenuation curve of Noll et al. (2009), and dust emission is included using the models of Draine & Li (2007). Moreover, a fraction of the stellar emission (f_{nodust}) is allowed to remain outside this dust screen. A complete description of the model parameters and adopted priors can be found in Wang et al. (2024b, Table 1) and Wang et al. (2025, Table 2).

The best-fit galaxy-only model implies a massive, post-starburst galaxy (Table 3) and a dust attenuation law that is extremely steep, with a dust index that reaches the edge of the prior (minimum dust index = -1.0). Only the combination of an evolved stellar population, with a low recent star formation rate (SFR), extremely low metallicity yet strong dust attenuation

Table 3. Stellar population fitting results.

Parameter	Prospector (galaxy-only/no AGN)	Prospector (galaxy + AGN)	Labbé et al. (galaxy + AGN)	Labbé et al. (galaxy + AGN; extreme dust)
$\log(M_*/M_\odot)$	$10.586^{+0.009}_{-0.048}$	$10.39^{+0.03}_{-0.03}$	$10.531^{+0.007}_{-0.008}$	$10.434^{+0.014}_{-0.049}$
age [Gyr]	$0.90^{+0.11}_{-0.42}$	$0.69^{+0.11}_{-0.11}$	$0.65^{+0.04}_{-0.04}$	$0.66^{+0.10}_{-0.19}$
$\log(Z/Z_\odot)$	$-1.95^{+0.14}_{-0.02}$	$-1.84^{+0.08}_{-0.09}$	$-0.986^{+0.023}_{-0.010}$	$-0.83^{+0.3}_{-0.12}$
SFR [$M_\odot \text{ yr}^{-1}$]	$7.3^{+6.0}_{-4.0}$	$0.43^{+0.95}_{-0.40}$		
$\log(\tau/\text{yr})$			$6.9^{+0.3}_{-0.3}$	$7.0^{+0.3}_{-0.4}$
A_V (stars)	$1.63^{+0.13}_{-0.08}$	$1.43^{+0.05}_{-0.05}$	$1.34^{+0.02}_{-0.02}$	$1.06^{+0.04}_{-0.03}$
dust index	$-0.991^{+0.013}_{-0.007}$	$-0.97^{+0.03}_{-0.02}$	-1.8	$-2.47^{+0.04}_{-0.02}$

Notes. The dust index corresponds the power-law slope for the Labbé et al. (2024) models. For the Prospector models it instead corresponds to the corrective power-law slope to the Calzetti dust law (i.e. δ in the parametrisation of Noll et al. 2009).

can produce a strong Balmer break as well as faint UV emission. However, Figure 5 (top panel) shows that this is still not sufficient to explain the observed spectrum, as significant residuals remain blueward and redward of the Balmer break. Because of the steep dust law, the model also slightly overpredicts the emission in the rest-frame near-IR spectrum and photometry (observed wavelength range of 4–8 μm).

5.1.2. Galaxy + AGN model

Next, we add an AGN model to perform a composite fit, following the procedure described in Wang et al. (2025). The intrinsic SED of this AGN accretion disc model is constructed from piece-wise power laws with indices fixed to the best-fit values of Temple et al. (2021). This SED is reddened by the same dust law that is applied to the stellar population, and additionally reddened by a second power law extinction curve of which the slope and normalisation are allowed to vary. We further include a model for the hot dust emission of an AGN torus, using the CLUMPY torus model (Nenkova et al. 2008) as implemented in FSPS (Conroy et al. 2009; Leja et al. 2018). The total AGN model therefore introduces five new free parameters: three of these parametrise the additional dust law (2) and optical depth of the torus (1), and the remaining two quantify the flux ratio between the galaxy and AGN at rest-frame 5500 \AA (f_{AGN}) and ratio of the AGN mid-IR luminosity and the galaxy bolometric luminosity ($f_{\text{AGN,torus}}$).

These fits can be tuned in order to down-weight or up-weight the galaxy versus AGN components, by setting different priors on f_{AGN} and the total mass formed. Wang et al. (2024a) hence proposed three different model flavours, resulting in models with a minimal, medium and maximal stellar contribution, with the maximal M_* model yielding the greatest Balmer break strengths. The priors used for the latter model are as described in Wang et al. (2025, Table 2), with the exception of f_{AGN} and the total mass formed, for which we assume a lognormal distribution with $\mu = -3$ and $\sigma = 1$ and a log-uniform distribution respectively.

Here, we present only this maximal M_* model in Table 3 and Figure 5, although we have also tested the other variants, finding significantly worse fits. The galaxy component dominates at the Balmer break and in the rest-optical, and the AGN only contributes substantially (although still sub-dominant) in the rest near-IR. The dust law required is similarly steep as for the galaxy-only model. Including the AGN component provides a marginally better fit in the rest-frame optical ($\sim 2\text{--}4 \mu\text{m}$) com-

pared to the galaxy-only model, but still cannot fit the Balmer break.

5.2. Labbé et al. composite model

Recently, Labbé et al. (2024) developed an independent full spectral fitting approach tailored to a luminous LRD with a strong Balmer break (the second strongest break in Figure 2). This fitting shares some of the features of the Prospector modelling, such as its Bayesian approach (performing nested sampling with pymultinest; Buchner et al. 2014), and the use of FSPS models with a Chabrier IMF. It differs in other aspects, such as the fact that only the spectrum (i.e. no photometry) is used for fitting. Moreover, the star formation history is parametrised as a delayed τ model, and the dust model is a power law with slope δ and variable normalisation. We use the exact same priors as described in detail in Table 6 of Labbé et al. (2024), with the exception of the dust law, for which we fix $\delta = -1.8$ (the lower bound of the prior range considered in their work). We assess the choice of this dust law in Section 6.1.

The most important differences are in the AGN component: this is modelled as a single power law continuum with variable slope β , and combined with a large suite of emission lines, such as the Balmer and Paschen series (with relative intrinsic fluxes assuming case B recombination), Fe II emission, and other lines commonly observed in AGN (e.g. He I, He II, Ne V, O I). Depending on line species, the emission is broadened according to the H α line width or the narrow [O III] emission. The AGN emission is reddened by a power law of the same slope as the dust law applied to the stellar population, but with a different normalisation.

Three different models were constructed with these components. The first model is a two-component blue+red AGN (where the blue AGN component is an unattenuated fraction of the red component), and the second a two-component galaxy model without any AGN. Neither model was shown to successfully produce a strong Balmer break. We therefore focus only on the third model here, which combines the two-component AGN model with a single stellar component. We make one modification to this model: because we do not observe any significant metal lines in *The Cliff* (Section 3.2), we only include broad hydrogen lines in this fit (neglecting He I).

The resulting fit is shown in the third panel of Figure 5. The best-fit AGN model is clearly substantially bluer than in the Prospector galaxy + AGN model. Otherwise, the fit looks remarkably similar to the two Prospector models, yielding a

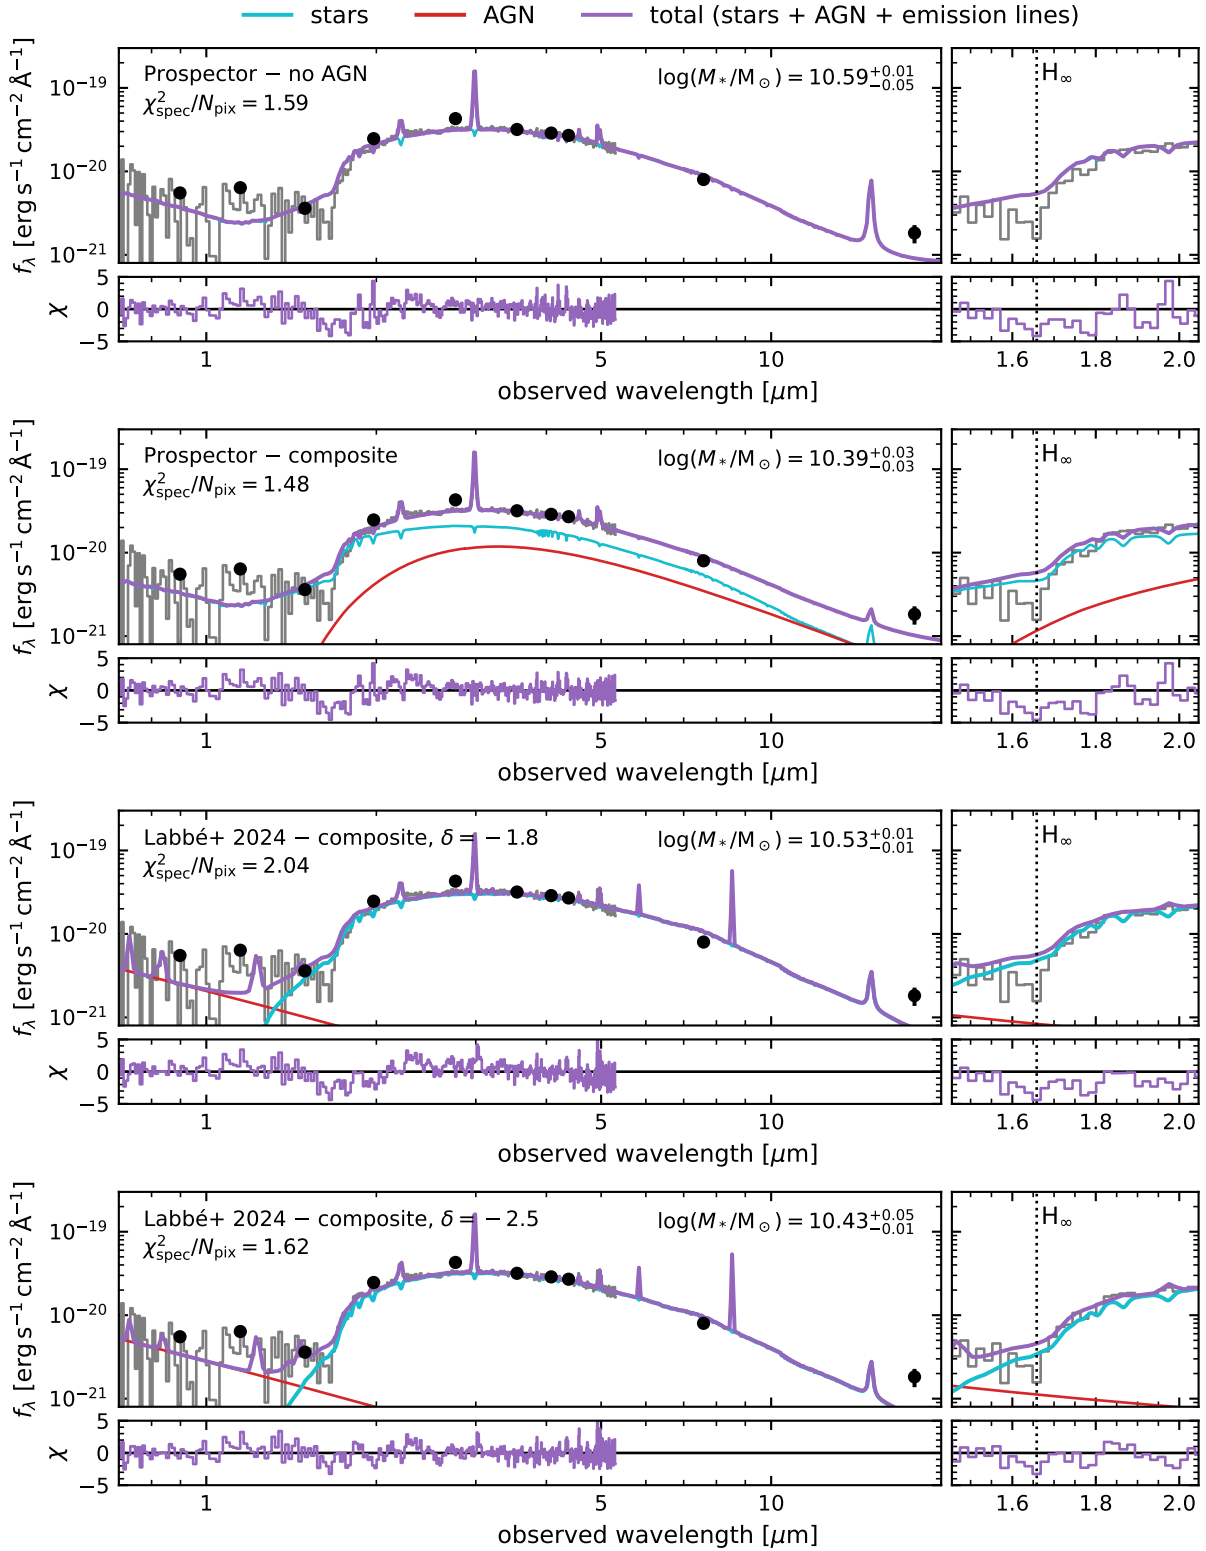


Fig. 5. Best-fit SED models and residuals (with respect to the PRISM spectrum) for four model variations, constructed as different mixtures of (dust-reddened) stellar population and power-law AGN model components. These AGN models do not consider the possible reddening by absorbing dense gas, which we explore in Section 6.4. The right panels show zoomed-in images of the region around the Balmer break. From top to bottom: the fiducial galaxy-only model from Prospector; a galaxy + AGN model that maximises the stellar contribution, fit with Prospector (the maximal M_* model of Wang et al. 2024a); the galaxy + AGN model following Labbé et al. (2024), but fitting only hydrogen emission lines instead of a forest of metal lines; the galaxy + AGN model of Labbé et al. (2024), but with an even steeper dust law (see Section 6.1). All four models favour a massive post-starburst solution, with a very steep dust attenuation law and high optical depth. However, even with such extreme dust, none of these models can produce the strong Balmer break and shape of the rest-frame optical SED, as is evident from the systematic (and significant) features in the residuals blueward and redward of the Balmer break.

massive post-starburst stellar population. As before, the combination of an older stellar population and a steep attenuation law is required in order to produce a strong Balmer break. The model fails in the same way as the Prospector models, as it does not capture the magnitude of the observed Balmer break, resulting in a systematic over- and under-prediction of the flux blueward and redward of the break, respectively. The model also slightly overpredicts the rest near-infrared, although the MIRI photometry was not included in the fit.

6. Discussion

We have shown that *The Clifff* has a highly unusual spectrum with a dramatic Balmer break that proves impossible to reproduce with typical evolved stellar populations. While massive post-starburst galaxies typically show the strongest Balmer breaks, as their SEDs are dominated by A stars, the exceptionally strong Balmer break of *The Clifff* vastly exceeds measurements of post-starburst galaxies at $z \gtrsim 3$. Only a steep dust attenuation curve can substantially increase this break strength. Therefore, if the rest-optical luminosity were dominated by starlight, the galaxy would have to be massive, strongly dust-attenuated, and abruptly ceased star formation in the last few hundred million years.

In what follows, we describe how we tested each of these characteristics. We first considered whether a dust law necessary to match the observed spectrum can be found, and whether it is realistic (Section 6.1). Next, we asked whether a high stellar mass within a very compact galaxy is plausible, by exploring the consequences of high stellar mass densities on the galaxy dynamics (Section 6.2). We then considered our assumption of a universal IMF and the contribution of A stars to the spectrum (Section 6.3), which underpins the SED modelling of Section 5. Finally, we discuss possible alternative models for *The Clifff*, in particular focusing on a new class of AGN-dominated models in Section 6.4.

6.1. Peculiar dust

The primary point here to emphasise about the SED models in Section 5 is that they imply an extreme, arguably implausible, combination of steep optical slopes and high optical depths ($A_V \sim 1.4$ – 1.7) for the dust attenuation. With power-law slopes $\delta \sim -1.8$, the implied UV-optical slopes $A_{1500}/A_V \sim 10$, a factor of ~ 2 steeper than any individual sightline in the SMC (Gordon et al. 2003). Most important in the context of this work is the reddening across the Balmer break, and instead of the UV-optical slope we therefore measure the optical slope of the dust law as the ratio of the attenuation at 3500 \AA and 5500 \AA (A_{3500}/A_V) in Figure 6. Clearly, both the galaxy-only model and the two composite AGN+galaxy models discussed in Section 5 have considerably steeper optical slopes (30–60%) than the Calzetti et al. (2000) dust law or the average Milky Way, SMC and LMC dust curves (Gordon et al. 2003, 2009, measured using the `dust_extinction` package of Gordon 2024).

Observations of large samples of galaxies (e.g. Salim et al. 2018), radiative transfer models (e.g. Chevallard et al. 2013, and references therein), and hydrodynamical simulations (at fixed grain composition, e.g. Narayanan et al. 2018) all show that such steep dust curves only occur in concert with low optical depths ($A_V \lesssim 0.5$), as the result of the star-dust geometry (for a review see Salim & Narayanan 2020). Conversely, high optical depths are observed and modelled only in systems with greyer attenua-




Fig. 6. Optical dust attenuation slope across the Balmer break (A_{3500}/A_V) vs. the optical attenuation A_V . Radiative transfer models predict a strong correlation between the slope and optical depth (Chevallard et al. 2013), which closely matches observed dust curves (Gordon et al. 2003, 2009; Salim et al. 2018). The various black symbols denote optical slopes and depths implied by the SED models of *The Clifff* from Section 5 (medians are shown with 16–84 percentiles, although these error bars are in most cases smaller than the markers), which deviate strongly from this trend. The dust-reddened BH* model is discussed in Section 6.4.

tion curves. The steep slopes and high optical depths of our SED models therefore lie outside the parameter space of observations and models so far (Figure 6), raising the question whether a suitable star-dust geometry or grain composition even exists that could produce such an attenuation curve.

Dust laws that are even steeper, yet optically thick at rest $\sim 4000 \text{ \AA}$, therefore appear physically implausible. Nevertheless, we ask if allowing for such a dust law would yield a better fit to the spectrum. We rerun our SED fitting for the Labbe et al. (2024) composite AGN+galaxy model with the power-law index as a free parameter, extending the lower bound of the prior on the dust law slope to $\delta = -2.5$. The UV-optical (optical) slope of this new bound is a factor of ~ 5 (~ 2) steeper than seen in SMC sightlines, and difficult to produce also in hydrodynamical simulations that investigate the dust-star geometry (Narayanan et al. 2018; Salim & Narayanan 2020). The posterior distribution for this model pushes against the boundary of $\delta = -2.5$ and with high $A_V \sim 1.2$, a factor of ~ 10 greater optical depth than expected for such a steep curve (Figure 6). The resulting model is shown in the bottom panel of Figure 5, and provides the best fit in the Balmer break region out of the four models presented.

However, significant residuals around the Balmer break remain, even with this extremely steep (and likely unrealistic) dust law. Previous studies of LRDs with strong Balmer breaks had already demonstrated that steep dust laws and high optical depths are necessary to fit the SED (Ma et al. 2025; Labbe et al. 2024; Wang et al. 2024a). But, the LRDs presented in these papers had Balmer breaks a factor of ~ 2 less strong than that of *The Clifff*, which in principle could be produced with the dust laws considered in Section 5.

Because the Balmer break of *The Clifff* is so exceptionally strong, we can now robustly conclude that massive stellar population models cannot fit the spectrum, even when considering extreme dust laws.

6.2. Kinematics of ultra-dense stellar systems

Before we assess a key assumption in our SED modelling – the shape of the IMF – we consider the implications that the

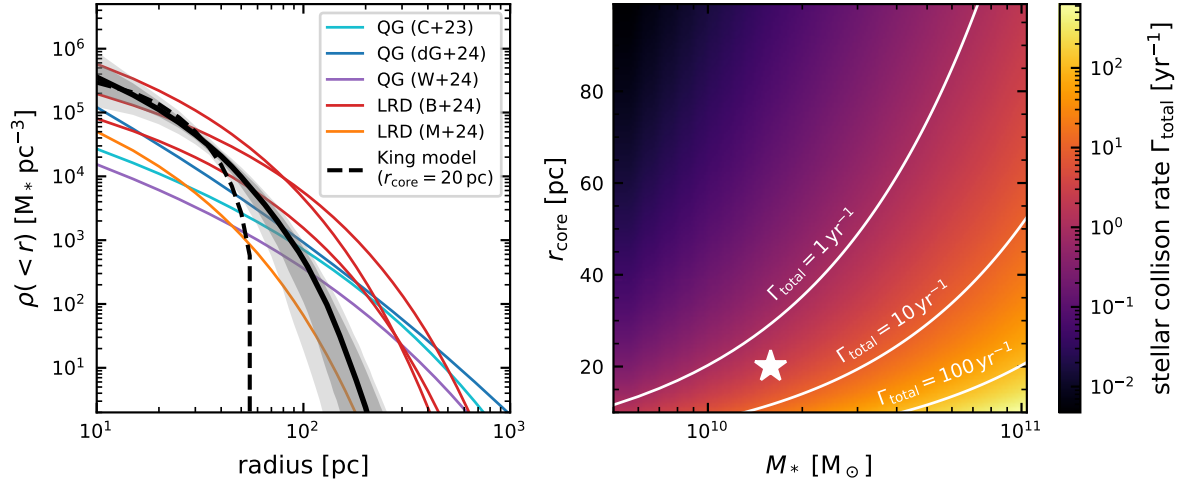


Fig. 7. Left: Stellar mass density profiles of *The Cliff* (black solid and grey shaded areas respectively indicate 16–84 and 5–95 percentiles), a selection of compact massive quiescent galaxies at $z > 4$ (cyan, blue, and purple lines; Carnall et al. 2023; de Graaff et al. 2025b; Weibel et al. 2025), and LRDs with strong Balmer breaks from the literature (orange and red lines; Baggen et al. 2024; Ma et al. 2025). The stellar mass densities of LRDs, under the assumption that the rest-optical is dominated by stars, exceed the densities in the cores of massive quiescent galaxies by an order of magnitude. Right: Stellar collision rate Γ_{tot} as a function of the core radius (as defined in the King model) and stellar mass of the system, for a ~ 500 Myr old stellar population. The inner $\lesssim 50$ pc of *The Cliff* can be well approximated with a King model of $r_{\text{core}} \sim 20$ pc; the location of the star corresponds to the dashed line shown in the left panel. If the stellar mass is as high as implied by the SED modelling of Section 5, we expect stellar collisions to occur at a frequency of $\sim 5 \text{ yr}^{-1}$ (a factor of $\sim 10^5$ greater than for the Milky Way’s nuclear star cluster).

high stellar masses of our explored SED fits would have on the galaxy’s dynamical properties. We found in Section 4 that *The Cliff* appears very compact ($r_{\text{e,maj}} \sim 40$ pc) and would have to have a stellar masses of $\log(M_*/M_\odot) \sim 10^{10.4-10.6}$ (Table 3), implying an extremely high mass density: the stellar mass surface density within the circularised effective radius $\Sigma_*(< r_e) \sim 6-9 \times 10^6 M_\odot \text{ pc}^{-2}$ exceeds the maximum value from observations and models of star clusters and galactic nuclei by over an order of magnitude ($\Sigma_{*,\text{max}} \sim 3 \times 10^5 M_\odot \text{ pc}^{-2}$; Hopkins et al. 2010; Grudić et al. 2019).

We deproject the Sérsic profile and estimate the 3D mass density profile, ρ , following Bezanson et al. (2009). This assumes spherical symmetry, and that the posterior distribution of Section 4 reflects an accurate measurement of the source morphology. The left panel of Figure 7 shows the median mass density profile (solid black line), with 16–84 (5–95) percentiles in dark (light) grey shading computed from 100 draws of the posteriors of the Sérsic profile fit and a range in stellar masses that bracket the different models of Table 3. We compare to the mass density profiles of massive quiescent galaxies at $z > 4$ (Carnall et al. 2023; de Graaff et al. 2025b; Weibel et al. 2025, cyan, blue, and purple lines), and also show the mass density profiles of LRDs with strong Balmer breaks, using stellar mass estimates that assume the rest-optical is dominated by stars (orange line; from Ma et al. 2025; Wang et al. 2024a) and size measurements (or upper limits) from Furtak et al. (2024) and Baggen et al. (2024, red lines). *The Cliff* is at least an order of magnitude denser than the cores of high-redshift massive quiescent galaxies, but matches the inferred (under the assumption that the optical is star-dominated) high stellar densities in some LRDs.

Such high stellar mass densities would imply that the stellar velocity dispersions are extremely high, with $\sigma_* \sim 10^3 \text{ km s}^{-1}$, as also discussed by Baggen et al. (2024). Guia et al. (2024) proposed that such high densities would in turn result in a system that is dynamically unstable to gravitational collapse, and Bellovary (2025) suggest that tidal disruption events from

the resulting runaway collapse could explain the properties of LRDs.

We consider both scenarios unlikely, at least for *The Cliff*. First, tidal disruption events that show broad Balmer emission lines typically have very blue SED shapes (reviewed in Gezari 2021), inconsistent with the observed SED of *The Cliff*. Second, the crossing time of a collisionless stellar system $t_{\text{cross}} = R/v \sim 10^{4-5} \text{ yr}$ for $R \sim 10^{1-2} \text{ pc}$ and $v \sim 10^3 \text{ km s}^{-1}$. Assuming that the system comprises $N = 10^{10}$ stars of equal mass, the timescale for dynamical relaxation (Eq. (1.38) of Binney & Tremaine 2008) is

$$t_{\text{relax}} = \frac{0.1N}{\ln N} t_{\text{cross}} \sim 10^{2-3} \text{ Gyr}, \quad (1)$$

and the core collapse timescale is approximately $t_{\text{cc}} = 0.2 \times t_{\text{relax}} \sim 10^{1-2} \text{ Gyr}$ for such dense stellar systems (Portegies Zwart & McMillan 2002), thus much longer than the age of the Universe at $z = 3.5$.

However, in the above we have treated stars as point-like particles. If the density is very high, and the system contains a significant fraction of evolved stars, we may expect stellar collisions to occur. In the following, we quantify the rate of such collisions, and explore whether there are observable consequences.

Generically, a local collision rate is defined via

$$\Gamma_{\text{local}}(r) = \frac{1}{2} n(r)^2 A \langle v_{\text{rel}} \rangle, \quad (2)$$

where $n(r)$ is the local number density [cm^{-3}], A is the collision cross-section [cm^2] = $X_*(M/M_\odot)\pi R_\odot^2$, X_* is the mean stellar disc area per unit mass in solar areas per solar mass, and v_{rel} is the relative velocity between stars. Integrating over the whole system leads to a total collision rate of

$$\Gamma_{\text{total}} = \int_0^{r_t} \Gamma_{\text{local}}(r) 4\pi r^2 dr.$$

Taking a King model for such dense stellar systems, one gets $n(r) = n_0 \exp(-\psi(r))$ for the density profile and

$$M = 4\pi m_* \int_0^{r_i} r^2 n(r) dr = 4\pi m_* n_0 r_c^3 I(\psi_i)$$

for the total mass, where $I(\psi_i)$ is a dimensionless function of the truncation parameter and m_* is the average stellar mass. The central velocity dispersion is related to the core radius: $\sigma_{v,0}^2 = 4\pi/9 G m_* n_0 r_{\text{core}}^2$. This eventually leads to a total collision rate of

$$\Gamma_{\text{total}}(\text{yr}^{-1}) \approx 3.7 \times 10^{-7} \alpha X_* \left(\frac{M}{10^5 M_\odot} \right)^2 \left(\frac{r_{\text{core}}}{1 \text{ pc}} \right)^{-5/2}, \quad (3)$$

for M in solar masses, r_{core} in parsecs, and X_* in solar areas per solar mass. The constant $\alpha \approx 0.1$ arises from the King profile (for a King parameter $W_0 \approx 7$).

Next, we need to calculate the cross section (to direct stellar collisions) per unit mass for a given stellar mass. For the vast velocities under consideration here, gravitational focussing does not dominate, and the cross section is simply geometric. We consider Padova isochrones (with a Kroupa IMF) and integrate over the full range of masses m : $X_* \equiv \int A_*(m) p_{\text{Kroupa}}(m) dm$, where $A_*(m) \equiv \pi R_*^2$ is the area of a star with mass m . For populations of several hundred Myr age we find values of $X \sim 5 \pi R_\odot^2 / M_\odot$. The right panel of Figure 7 shows the total stellar collision rate as function of mass and core radius. For *The Cliff*, a King model with $r_{\text{core}} = 20 \text{ pc}$, $M = 10^{10.3-10.4} M_\odot$ (i.e. approximately half the stellar mass implied by the SED models) and $\alpha = -0.1$ provides a good description of the inner density profile (black dashed line); for $X = 5$ this leads to $\Gamma_{\text{total}} \sim 4-6 \text{ yr}^{-1}$. Even for the lowest stellar mass found in Section 5 and upper limit on the size, we still find $\Gamma_{\text{total}} = 1 \text{ yr}^{-1}$. Previously, such high-speed collisions have only been considered for the immediate vicinity of supermassive black holes (creating velocities of $\gtrsim 1000 \text{ km s}^{-1}$; Ryu et al. 2024; Hu & Loeb 2024), where the collision rates are dramatically lower (e.g. $10^{-4.5} \text{ yr}^{-1}$, for the Milky Way’s nuclear cluster).

What happens if the envelope of an evolved star is pierced in a random spot, e.g. at $R_*/\sqrt{2}$, by a solar-type main sequence star deserves, and quantitatively requires, hydrodynamical simulations that are beyond the current scope. High-speed stellar collisions have been simulated hydrodynamically in the literature (e.g. Ryu et al. 2024), but not in the regime most likely here. Instead, we can offer an order-of-magnitude toy scenario. The main sequence star, acting at $\sim 800 \text{ km s}^{-1}$ as a ‘bullet’ piercing the evolved star’s envelope, will create a bow shock that exceeds in extent the radius of the piercing star by a few times (3–5 \times). The material then swept up within a cylinder of $5 R_\odot$ at a projected centre distance of $40 R_\odot$ in an evolved star of $R \approx 60 R_\odot$ and $3 M_\odot$ will sweep about $0.01 M_\odot$ of the material. One can then envision that this material is heated to the virial temperature of the collision, of about $5 \times 10^7 \text{ K}$ for a collision at 800 km s^{-1} , corresponding to a thermal energy of $10^{47.5} \text{ erg}$ and a blackbody peak in the X-rays at 4.5 keV . How much of this energy will be radiated away, and how much will drive the expansion is unclear. The thermal energy alone could afford a luminosity of $10^{42.5} \text{ erg s}^{-1}$ for a week, which is shorter than the time between collisions and therefore might appear as X-ray outbursts. Although the expected luminosity is above the existing X-ray flux limits for *The Cliff* (Section 2.3), we cannot definitively rule out stellar collisions in this particular system due to the likely time variable nature of the X-ray emission.

These considerations leave open a number of fundamental issues in considering such ultra-dense systems. Foremost, the questions arise of how such systems could ever have formed from an ISM and how stellar evolution proceeds in an environment of frequent stellar envelope piercing.

6.3. Variations in the initial mass function

Thus far, we have made the assumption that the IMF of the system resembles that of the Milky Way. At higher redshifts, however, this assumption may no longer hold due to different conditions of the interstellar medium. Top-heavy IMFs have recently been explored as a solution to explain the discovery of extremely massive and extremely UV-luminous galaxies at $z \gtrsim 6$ (e.g. Harikane et al. 2023b; Steinhardt et al. 2023; Woodrum et al. 2024; Cueto et al. 2024; Mauerhofer et al. 2025). Such IMFs increase the ratio of massive to dwarf stars and therefore can reduce the mass-to-light ratio (M/L). In the context of this work, top-heavy IMFs are interesting because a higher fraction of stars $> 1 M_\odot$ could also increase the Balmer break strength.

Exploring variations in the IMF in a self-consistent modelling framework such as Prospector is challenging, and the parameter space to explore (i.e. IMF parametrisation, star formation history, dust law) is potentially extremely large. We therefore take an empirical approach instead, and leverage the fact that, ultimately, any stellar population synthesis model represents a weighted average of a large and diverse set of stellar spectra. Moreover, because *The Cliff* shows an extremely strong Balmer break and its blue rest-frame near-infrared SED is suggestive of the Rayleigh-Jeans tail of a black body, it is reasonable to assume that a single component (i.e. a narrow range in stellar effective temperatures) dominates such a weighted average.

We therefore search for stars that most closely resemble the SED of *The Cliff* in one of the most commonly used stellar libraries for constructing SPS models, MILES (Sánchez-Blázquez et al. 2006), which was also used in our modelling in Section 5. To narrow down our search, we compute Balmer break strengths for the spectra using the same definition as in Section 3 and further inspect only those stars with strong breaks. The distribution of break strengths is shown in the top left panel of Figure 8 (blue). Although the shape of this distribution is arbitrary, i.e. it does not reflect our assumed IMF, this reveals that the maximum possible break strength in the stellar library ~ 4 , significantly lower than the measured Balmer break of *The Cliff* (black solid line). If we apply a Calzetti et al. (2000) dust law with $A_V = 1$ (dashed purple), the break strengths increase slightly, but there are still no stars that match the observed Balmer break. We also apply the steep dust attenuation of the galaxy-only Prospector model from Section 5.1 (red solid), which increases the typical Balmer break by a factor of ~ 2 and yields a subset of stars with break strengths comparable to *The Cliff*.

Next, we select the strongly attenuated stars with Balmer breaks > 5.5 , which overlap with the 1σ uncertainty of *The Cliff*. The distribution of effective temperatures (top right) is bimodal: the majority of stars have $T_{\text{eff}} \sim 4000 \text{ K}$, and a smaller subset is clustered at $T_{\text{eff}} \sim 8000 \text{ K}$. Upon visual inspection, we identify four types of stars and show a representative spectrum for each type in the bottom panel of Figure 8 (where we have convolved the MILES spectra to the PRISM resolution at $z = 3.5$). The stars with low T_{eff} have the strongest breaks (up to a value of ~ 8), and are giant stars with strong metal absorption features blueward of 4000 Å , leading to the well-known 4000 Å break. Although these giant stars do have strong Balmer breaks by the

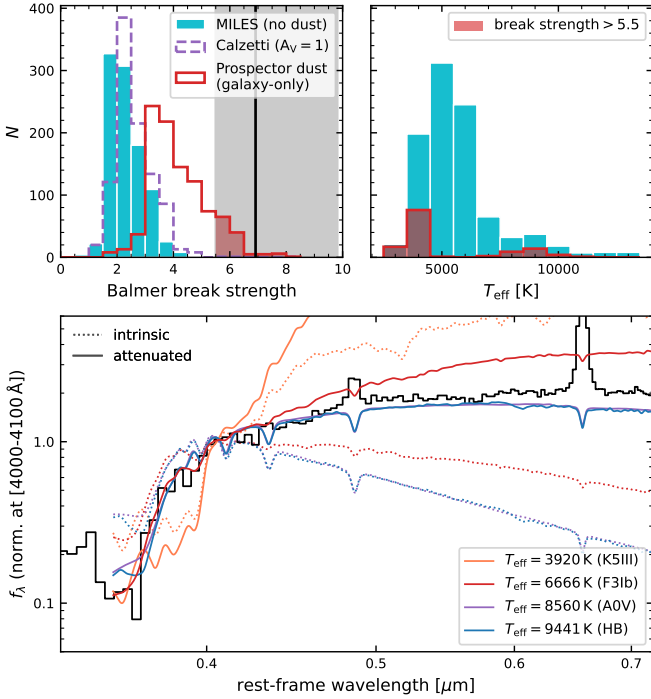


Fig. 8. Top left: Balmer break strength (from rest-frame [3620, 3720] Å and [4000, 4100] Å ranges) distribution for the standard MILES stellar library (blue), and after applying a Calzetti dust law (with $A_V = 1$; dashed purple), and the steep attenuation law from the Prospector galaxy-only model (solid red). The black vertical line and shaded region show the measured break strength and 1σ uncertainty of *The Cliff*. Top right: Effective temperatures of all MILES stars (blue) and the stars that, after applying the steep dust law, fall within 1σ of *The Cliff* (red). Bottom: Representative sample of MILES spectra with strong Balmer breaks, convolved to the PRISM resolution at $z = 3.5$ and normalised at [4000, 4100] Å. Dotted and solid lines show the stellar spectra without dust and with steep dust attenuation, respectively. In comparison to the spectrum of *The Cliff*, K giants, main-sequence A stars and horizontal branch stars fall short at <3645 Å and >5000 Å. Only massive supergiants with steep dust attenuation can match the strength and shape of the Balmer break (albeit with a mismatch at >5000 Å), and suggests that an extraordinary, top-heavy IMF would be required to explain the rest-frame optical emission of *The Cliff*.

(broad) definition we have used, their spectral shapes clearly do not match the SED of *The Cliff*.

At $T_{\text{eff}} \sim 8000\text{--}10000$ K main sequence A stars yield Balmer break strengths ≤ 6.5 after applying steep dust attenuation; horizontal branch stars in this temperature range show marginally stronger breaks. In detail, however, Figure 8 shows that these stars are not sufficiently blue at <3645 Å and not red enough at >5000 Å. This is similar to our findings in the SED modelling of Section 5. Importantly, it implies that, even if A stars were to dominate the SED entirely, *The Cliff* cannot be modelled as a normal post-starburst galaxy.

Intriguingly, we identify one star in the MILES library that matches the Balmer break of *The Cliff* remarkably well (after substantial dust attenuation; red line in Figure 8), although the two spectra start to diverge at >4500 Å. Its spectral type, F3 Ib, corresponds to a yellow supergiant with $T_{\text{eff}} \sim 6700$ K. Such stars represent a brief phase ($\sim 10^{-4}$ yr) in the evolution of a massive star (initial mass $\sim 8\text{--}20 M_{\odot}$; e.g. Drout et al. 2009).

This suggests a stellar population with an extraordinarily top-heavy IMF – coupled with a very steep dust law – may be

able to explain the SED of *The Cliff*. It would, however, require a great degree of fine tuning of the star formation history to achieve the right composition of A and supergiant stars needed to match the spectrum, as the supergiant phase is extremely short-lived and the presence of any O or B stars would wash out the strong Balmer break. Importantly, such an IMF would drastically lower the M/L and hence also reduce the stellar mass density. We estimate the change in M/L (at 5500 Å) for extreme IMFs by constructing simple stellar populations for a range of δ -function IMFs in FSPS, attenuated by the same dust curve as the Prospector galaxy-only model. To do so, the IMF is parametrised by two power laws with a steep negative slope (γ_1) for masses $m < m_0$ and positive slope (γ_2) for $m > m_0$, with transition mass $m_0 = 2 M_{\odot}$ (i.e. main-sequence A stars). For stellar populations dominated entirely by A stars ($\gamma_1 = -50$, $\gamma_2 = 50$, ages $\sim 0.1\text{--}1$ Gyr), the M/L is a factor of ~ 10 lower than for the models of Section 5. If we allow for a larger fraction of massive stars (by varying $\gamma_2 \sim 1\text{--}2$), we find differences in M/L of a factor of $\sim 10\text{--}50$ depending on the age. The lower mass density of the system in turns implies a considerably lower stellar velocity dispersion, by a factor of $\sim 3\text{--}7$ ($\sigma_* \sim 100\text{--}200$ km s $^{-1}$).

Recent work by van Dokkum & Conroy (2024) proposed that a so-called ski-slope IMF (i.e. a steep slope at $<0.5 M_{\odot}$ and shallow slope at $>1 M_{\odot}$) could reconcile observations of massive galaxies at high and low redshifts, as such an IMF would look top-heavy at high redshifts due to the higher fraction of massive stars, while a large population of low-mass stars would appear as a bottom-heavy IMF in observations of local early-type galaxies. It may be possible to conceal a large population of low-mass stars ($<0.5 M_{\odot}$) in *The Cliff*, but the increased number density of both dwarf and giant stars in such a scenario would also substantially increase the cross section of the stellar population. Although a detailed calculation, with revised stellar mass and hence velocity dispersion estimates, is necessary to determine the corresponding change in the stellar collision rate and X-ray luminosity, it is likely that such a system would violate current X-ray constraints.

We therefore conclude that it may be theoretically possible to construct a star formation history, dust law, and top-heavy IMF that reproduces the optical to near-IR continuum of *The Cliff* (at a lower stellar mass density), but that the SED resulting from such a contrived setup would be short-lived and therefore extremely rare. Yet, the fact that a second source with near-identical SED shape and luminosity has been found in the same small survey area of ~ 200 arcmin 2 (Figure 9; Naidu et al. 2025) appears to argue against these sources being rare transient phenomena. In the following section, we therefore proceed to discuss alternative models that do not require any prominent stellar component in the rest-optical.

6.4. A black hole star

In Section 3.2 we showed that the H α emission complex of *The Cliff* contains a broad Lorentzian component of FWHM ~ 1500 km s $^{-1}$. Although we do not detect any forbidden lines (e.g. [O III]) and thus lack a strong constraint on the narrow line kinematics, such a broad profile points to the likely presence of an accreting black hole. Assuming the single epoch scaling relation between black hole mass, H α luminosity and H α line width of Greene & Ho (2005), we obtain a black hole mass of $\log(M_{\text{BH}}/M_{\odot}) = 7.18^{+0.07}_{-0.06}$, as we will discuss further below, we do not include a reddening correction, but note here that M_{BH} increases by 0.54 dex for $A_{\text{H}\alpha} = 2$. The bolometric luminosity implied by the H α line luminosity (assuming the scaling

relations between $L_{H\alpha}$, L_{5100} , and L_{bol} of Greene & Ho 2005 and Shen et al. 2020), $L_{bol} \sim 1 \times 10^{45} \text{ erg s}^{-1}$, is a factor of ~ 5 – 10 higher than the bolometric luminosity implied by the 3σ upper limit on the soft X-ray luminosity, but consistent with the 3σ upper limit from hard X-rays (see Section 2.3).

The high observed equivalent width of the broad $H\alpha$ line, $EW_{H\alpha,rest} \sim 400 \text{ \AA}$, suggests that AGN emission must also account for a substantial fraction of the rest-optical continuum. However, in Section 5 we showed that typical AGN models, i.e. blue SEDs described by piece-wise power laws, cannot reproduce the spectrum of *The Clifff* even after strong dust attenuation. The difficulty in finding a suitable AGN model is that it needs to fit the Balmer break region, produce a red rest-optical continuum, yet a comparatively blue rest-frame near-IR. The absence of X-ray emission as well as hot dust emission in the mid-IR (at least out to rest $4 \mu\text{m}$) pose additional challenges.

Interestingly, our exploration of the MILES stellar library in the previous section showed that some stellar spectra (to first order) are a good match to *The Clifff*. Inayoshi & Maiolino (2025) recently proposed that dense gas clouds in the vicinity of an AGN accretion disc could mimic some of the conditions of a stellar atmosphere, such as the high gas density ($n_H \sim 10^9 \text{ cm}^{-3}$), resulting in a strong Balmer break and Balmer absorption lines. Ji et al. (2025) implemented such a model in the photoionisation software Cloudy (Ferland et al. 2017) and showed that an incident AGN spectrum transmitted through a slab of dense gas with high turbulent velocity can indeed produce both a strong Balmer break and strong Balmer absorption. They hence showed that, after applying an SMC attenuation law with $A_V = 2.1$, this results in a reasonable fit to the spectrum of the triply imaged LRD A2744-QSO1 of Furtak et al. (2024), with a Balmer break that is slightly stronger than observed. Although the Balmer break of A2744-QSO1 (shown in Figure 1, red line) is less strong (by a factor of 2) than that of *The Clifff*, the two sources do share some other characteristics, such as luminous and broad Balmer emission, very weak [O III] emission, and narrow Balmer absorption.

6.4.1. Dust-reddened BH* models struggle in the infrared

Such a model, hereafter referred to as a black hole star (BH*), may therefore also be a good description of the extremely strong Balmer break and absorption in the $H\alpha$ line seen in *The Clifff*. Although this term has some similarities to the quasi-stars proposed in the past to explain the rapid growth of light seeds into intermediate-mass black holes (e.g. Begelman et al. 2006, 2008), the BH* model does not necessarily imply spherical symmetry, hydrostatic equilibrium, or require the collapse of pristine gas. During the writing of this paper Naidu et al. (2025, hereafter N25) discovered a source (named MoM-BH*-1) with a near-identical spectral shape to *The Clifff* in the same imaging area of the UDS, but at $z_{spec} = 7.76$. We show in Figure 9 that these sources differ only by a factor of ~ 2 in luminosity, although the Balmer emission of MoM-BH*-1 is a factor of ~ 2 broader in width. As discussed in detail in N25, the spectrum of MoM-BH*-1 can be well described with an AGN-dominated model in which a shell of dense gas (i.e. with a high covering fraction, for simplicity set to 100%) surrounds the accretion disc of a massive black hole. In what follows, we will use the modelling framework presented in their work. Briefly, the free parameters of this model consist of the gas density (n_H), column density (N_H), metallicity ($\log Z$), ionisation parameter ($\log U$), and the turbulent velocity (v_{turb}); the AGN accretion disc spectrum uses

the default AGN models of Cloudy, parametrised by the temperature of the big blue bump and three power laws. The transmitted spectrum is then computed with Cloudy, where in practice the 3D spherical shell model is simplified to a slab of gas, and both the continuum and line emission from the dense gas itself are taken into account. Lastly, we note that, because only a finite number of hydrogen energy levels are computed, there is an artificial emission line at $\sim 3650 \text{ \AA}$ from the blending of higher-order transitions.

The parameter space of these BH* models is high-dimensional, and directly fitting for the best model is at present not feasible. We therefore start by examining a BH* model that is broadly similar to the one proposed by Ji et al. (2025), with $n_H = 10^{10} \text{ cm}^{-3}$, $N_H = 10^{24} \text{ cm}^{-2}$, $\log(Z/Z_\odot) = -1.5$ and $\log U = -1.5$. A high turbulent velocity ($v_{turb} = 500 \text{ km s}^{-1}$) is needed to produce the EWs of the Balmer absorption features seen in MoM-BH*-1. Whether this is also needed for *The Clifff* is unclear, primarily due to the poorly constrained absorption feature in the G395M spectrum (Section 3.2). For the purpose of this paper we therefore focus only on the shape of the continuum and defer a detailed modelling of the emission lines and Balmer absorption to a future study with deeper G395M data (GO-7488; PI: A. Weibel). The incident AGN spectrum is described by $T = 10^5 \text{ K}$ and $\alpha_X = -1.5$, $\alpha_{UV} = -0.5$, $\alpha_{ox} = -1.5$ for the X-ray, UV and X-ray to optical power-law slopes, respectively. This model was drawn from a large grid ($\sim 10^6$ models) and chosen to produce a qualitatively similar spectral shape (Balmer break strength, UV and optical continuum slope) as Ji et al. (2025). The model is normalised to the spectrum of *The Clifff* by the luminosity at rest $[4000, 4100] \text{ \AA}$ and convolved to the PRISM resolution at $z = 3.55$.

The intrinsic BH* model (top panel of Figure 9; dark green line) has a Balmer break strength ~ 4 , stronger than that of any star with $T_{eff} > 6000 \text{ K}$ explored in Section 6.3, and therefore performs better than the stellar population models attempted before. However, the model slightly overpredicts the rest-UV emission, and severely underpredicts the emission in the rest-optical and near-IR. We next apply an average SMC dust law with the low optical depth observed in the SMC ($A_V = 0.5$; Gordon et al. 2003). This model (top panel, green line) provides a remarkably good fit in the rest-frame UV and up to rest $\sim 4200 \text{ \AA}$, but falls short at longer wavelengths. Following Ji et al. (2025) we also show the SMC attenuation with $A_V = 2$, although we caution that this dust law suffers from exactly the same problems discussed in Section 6.1, as the steepening of the SMC dust law has been shown to arise from a specific star-dust geometry at low optical depth (e.g. see Chevallard et al. 2013; Salim & Narayanan 2020). The light green line shows this attenuated model, providing an excellent match to the region around the Balmer break (~ 3645 – 4800 \AA), and a reasonable fit out to $H\alpha$. The rest-UV emission is nearly completely absorbed, although this could be remedied, for example, by invoking a low-mass star-forming host galaxy outside of the hydrogen envelope or allowing for scattered light from the AGN by lowering the covering fraction.

More importantly, this strongly dust-reddened BH* model reveals a critical issue in the rest near-IR, as the flux at wavelengths beyond $H\alpha$ deviates significantly from the observed spectrum. This issue is not apparent in either of the NIRSpec spectra modelled by Ji et al. (2025) and N25, as both their objects are at $z \sim 7$ – 8 where NIRSpec provides coverage out to the rest-optical and only MIRI data is able to constrain the rest IR. With complete coverage from X-ray wavelengths to the

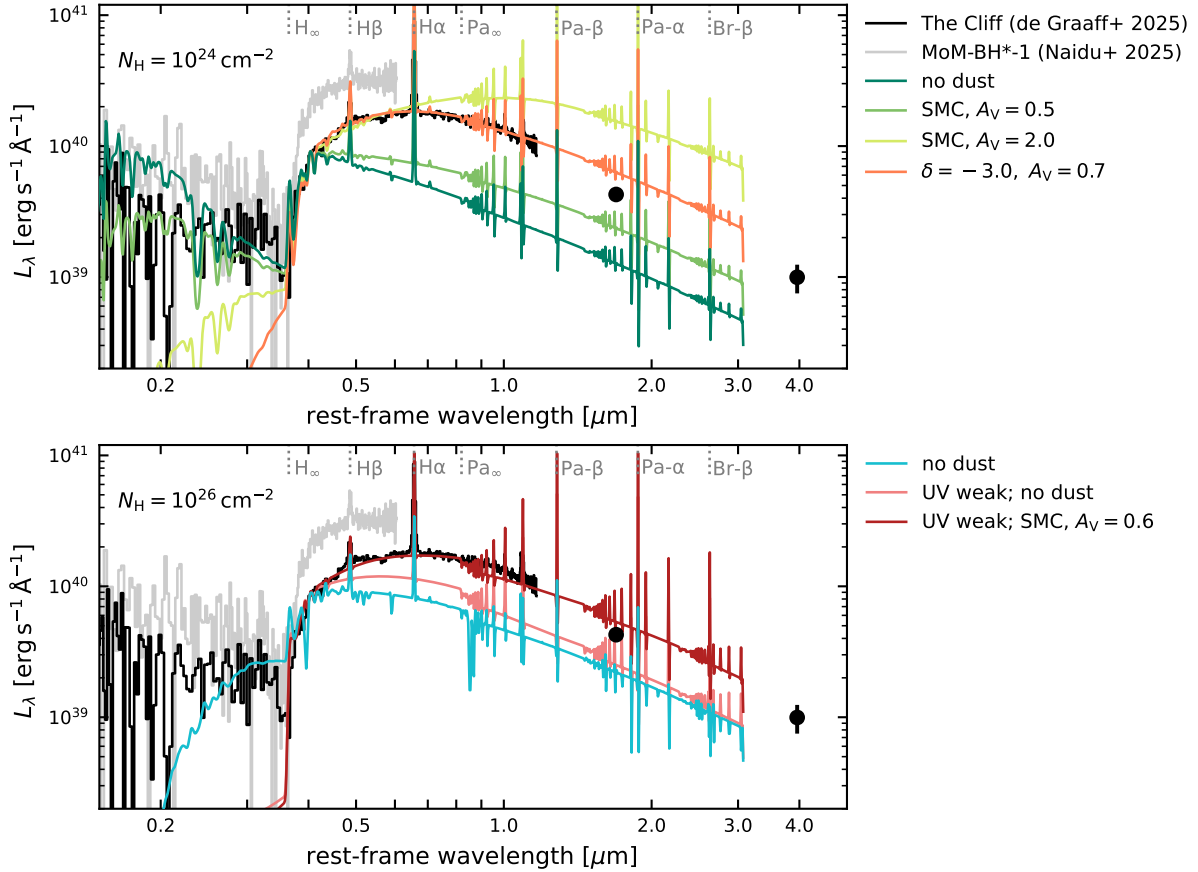


Fig. 9. NIRSpect/PRISM spectra of *The Cliff* and MoM-BH*-1 ($z_{\text{spec}} = 7.76$; Naidu et al. 2025). They are remarkably similar in shape, and only differ by a factor of ~ 2 in luminosity. Top: BH* star model of a blue incident AGN spectrum ($\alpha_X = -1.5$, $\alpha_{\text{UV}} = -0.5$) with high column density of $N_{\text{H}} = 10^{24} \text{ cm}^{-2}$ (rescaled to the rest [4000, 4100] Å flux of *The Cliff* and convolved to PRISM resolution) yields an intrinsically strong Balmer break (dark green) that matches the spectrum better than any of the stellar spectra explored (Figure 8). Steep dust curves are required to match the shape of the spectrum blueward of ~ 4200 Å (green), and high optical depths are needed to be able to match the curvature of the rest-optical spectrum (light green, orange). However, such strongly dust-reddened AGN models result in a severe discrepancy in the rest near-IR. Bottom: Increased absorption from dense gas ($N_{\text{H}} = 10^{26} \text{ cm}^{-2}$) results in extra reddening of the spectra without the need for dust (blue line), but the Balmer break of this model does not match the spectrum. A UV-weak incident AGN spectrum ($\alpha_X = -0.5$, $\alpha_{\text{UV}} = -0.1$) can produce a strong Balmer break as well as a redder rest near-IR, although small mismatches in the IR remain even after modest dust attenuation (red lines). This suggests a redder intrinsic AGN spectrum (as predicted by some super-Eddington accretion disc models), or additional component in the SED such as (super)massive stars, is needed to explain *The Cliff*.

rest mid-IR, *The Cliff* therefore provides a unique test-bed for future BH* models.

6.4.2. Alternatives to dust reddening

The question arises of what is needed in order to fit the SED of *The Cliff*. The fact that the BH* model is able to fit the rest-UV and Balmer break (up to ~ 4200 Å) with minimal dust attenuation is encouraging, and suggests that this class of models may provide a reasonable description of the underlying physics. However, the discrepancy in the rest-optical and near-IR clearly indicates that some modification is necessary. It is possible to achieve extra reddening in the near-IR with a dust law steeper than that of the SMC, as shown in the top panel of Figure 9 (orange line). But, invoking such a dust law with high optical depth is problematic (see Section 6.1), and still cannot produce a spectrum that simultaneously matches the rest-optical and near-IR. It may also be difficult to reconcile strong dust attenuation with the non-detection of hot dust in the mid-IR, and we find that *The Cliff* is also not detected in archival Spitzer/MIPS and Herschel/PACS data (obtained from S. McNulty

et al., in prep.), although deep far-IR observations would be needed to robustly constrain the presence of substantial warm or cold dust. Setton et al. (2025) report simultaneous mid- and far-IR observations for two different luminous LRDs and show that the maximum IR luminosity limits are in tension with models that invoke strong dust attenuation. Similarly, deep far-IR observations of (bright) LRDs at $z \sim 3\text{--}7$ thus far by Xiao et al. (2025) and Akins et al. (2025b) reveal no (cold) dust continuum emission. Only Barro et al. (2024b) report a compact red source that is well-detected in the mid- and far-IR indicative of hot dust, although this source does not yet have a spectroscopically confirmed v-shaped rest-UV to optical continuum or broad Balmer line, and it is therefore unclear whether this is a LRD. Under the assumption of energy balance, the BH* model attenuated by an SMC dust law with $A_V = 2.0$ would yield an IR luminosity of $10^{11.2} L_{\odot}$, which is only 0.5 dex below the 3σ limit on the total IR luminosity for LRDs measured in Setton et al. (2025, after normalising to rest 6000 Å), and thus detectable with ground-based sub-mm facilities.

We therefore propose that the AGN SED may be intrinsically red (at rest-optical wavelengths), rather than dust-reddened. This

could be achieved by increased reddening from dense absorbing gas, i.e. by increasing the gas density. The bottom panel of Figure 9 shows the same BH* model as before, but for a higher density of $n_{\text{H}} = 10^{11} \text{ cm}^{-2}$ and $N_{\text{H}} = 10^{26} \text{ cm}^{-2}$ (blue line). This results in a redder rest-optical and rest near-IR continuum, but the resulting Balmer break is not sufficiently strong. Alternatively or in addition, the accretion disc spectrum could be redder than that of typical AGN, which has been predicted by some super-Eddington models (e.g. Abramowicz et al. 1980; Jiang et al. 2014, 2019) and, separately, models of accreting black holes inside dense gaseous envelopes (i.e. quasi-stars) have been linked to super-Eddington accretion (albeit for lower black hole masses, e.g. Volonteri & Begelman 2010; Coughlin & Begelman 2024). Moreover, super-Eddington accretion has also been suggested as a solution to the observed X-ray weakness of LRDs (Lambrides et al. 2024; Yue et al. 2024; Ananna et al. 2024; Inayoshi et al. 2024), although the presence of extremely dense ($N_{\text{H}} = 10^{26} \text{ cm}^{-2}$), Compton-thick gas in the BH* model may alone be sufficient to explain such X-ray weakness. Using the bolometric luminosity and black hole mass estimated from the broad $\text{H}\alpha$ line under the assumption that local scaling relations apply to *The Cliff*, we obtain an Eddington ratio $L_{\text{bol}}/L_{\text{edd}} \approx 1.3$. However, as no source similar to *The Cliff* has been discovered before, it is unclear that commonly used scaling relations can be applied. The broad wings of the Lorentzian line profile may be due to other processes, such as resonant scattering (see N25) or electron scattering (Laor 2006; Rusakov et al. 2025), and would imply that the broad line width does not reflect the kinematics of the broad line region. The black hole mass may therefore be severely overestimated (by up to ~ 2 dex), which in turn would strongly increase the Eddington ratio ($L_{\text{bol}}/L_{\text{edd}} \sim 10$; N25).

Two models are shown in the bottom panel of Figure 9 (red lines) for extremely very high column densities ($N_{\text{H}} = 10^{26} \text{ cm}^{-2}$) and AGN spectra that have shallower X-ray and UV slopes ($\alpha_{\text{X}} = -0.5$, $\alpha_{\text{UV}} = -0.1$, $\alpha_{\text{ox}} = -1.5$), corresponding to the fiducial model parameters of N25. These intrinsically redder models provide a better match to the overall shape of the rest-optical and near-IR continuum. Nevertheless, small yet significant mismatches in the near-IR (rest $\sim 0.8\text{--}1.2 \mu\text{m}$) remain, even after modest dust attenuation (SMC law with $A_{\text{V}} = 0.6$). In our initial exploration of the large parameter space spanned by the AGN and gas properties, we tested only a limited variety of AGN models (spanning a grid of 10^6 models). It is therefore possible that a better model does exist in a larger parameter grid, but we defer such a detailed empirical investigation of the intrinsic AGN SED to a future paper.

The fact that none of the BH* models presented above can match the spectrum of *The Cliff* in detail could also imply that we are missing a key ingredient, such as a stellar component. For instance, AGN models with intrinsically red SEDs have been proposed in the past which do not require super-Eddington accretion rates, but instead impose a heating of the outer accretion disc by (super)massive stars (e.g. Sirko & Goodman 2003, although these models may be too luminous in the mid-IR). More generally, we have so far assumed that an accretion disc powers the BH*, but in principle any powerful ionising source inside the dense gaseous envelope could produce a strong Balmer break and luminous emission lines.

Nuclear star clusters, possibly surrounding a massive black hole, may contain such a large population of massive stars with hard ionising spectra: observations of the nuclear star cluster in the Milky Way have revealed an extremely top-heavy IMF (Bartko et al. 2010; Lu et al. 2013). If such a nuclear star

cluster also contains supermassive stars, which are expected to have modest effective temperatures and hence red spectra (e.g. Martins et al. 2020), this may also help to explain the rest near- and mid-IR SED of *The Cliff*. Interestingly, the formation of nuclear star clusters and the growth of massive black holes have been demonstrated to be tightly linked, and runaway collisions in nuclear star clusters have long been proposed as a formation channel of massive black holes (for a review, see Neumayer et al. 2020). Goodman (2003) showed that massive stars could even form in the outer accretion disc of a massive black hole, and therefore could be encompassed in the same dense gaseous envelope. To what extent the stellar emission would contribute to the SED, and whether this scenario can produce the broad symmetric Balmer emission lines is currently unclear. Moreover, dedicated models are needed to assess the effects of stellar feedback on the dense gas in such a system and determine whether this would be a short-lived phase or a stable mode of star formation. Considering the high number density of LRDs (e.g. Kokorev et al. 2024), a predominantly AGN origin of the rest-optical to near-IR SED of *The Cliff* therefore appears most plausible at face value, but stellar emission cannot yet be ruled out, and may even contribute substantially.

7. Conclusions

We present a detailed investigation of *The Cliff*, a luminous LRD at $z_{\text{spec}} = 3.55$ with an extremely strong Balmer break discovered in RUBIES. With high-quality JWST/NIRSpec PRISM and G395M spectroscopy from RUBIES, JWST/NIRCam and MIRI imaging from the PRIMER survey, and archival Chandra X-ray data, it is one of the few (bright) LRDs with complete spectrophotometric coverage from X-ray to mid-IR wavelengths. Coupled with its highly unusual SED, *The Cliff* provides a crucial stress test of stellar population and AGN models.

The spectral and morphological properties of *The Cliff* present a complex puzzle:

- The SED is characterised by a Balmer break that exceeds the break strengths of $z > 3$ massive quiescent galaxies and LRDs published thus far by a factor of 2. Together with a (weak) blue UV slope, this gives rise to the v-shape typically observed in LRDs. The PRISM spectrum (extending to rest $1.2 \mu\text{m}$) and MIRI photometry (rest $\sim 2\text{--}4 \mu\text{m}$) reveal a blue SED in the near- and mid-IR, indicating there is no evidence of hot dust emission at least to rest $\lesssim 4 \mu\text{m}$.
- The PRISM spectrum reveals a suite of Balmer, Paschen, and He I emission lines, but no significant metal lines, pointing to low-metallicity or high-density gas. The G395M spectrum kinematically resolves $\text{H}\alpha$ and various Pa lines, showing broad Lorentzian profiles ($\text{FWHM} \sim 1500 \text{ km s}^{-1}$) as well as a narrow redshifted absorption feature in the $\text{H}\alpha$ line. If we assume this originates from the broad-line region of an AGN and use low-redshift single-epoch black hole mass scaling relations, the $\text{H}\alpha$ line translates to a black hole mass $M_{\text{BH}} \sim 1 \times 10^7 M_{\odot}$.
- Sérsic profile fitting to NIRCam F200W imaging shows that the source is extremely compact, but possibly marginally resolved compared to a nearby star. The single-component Sérsic model converges to $r_{\text{e,maj}} \sim 40 \text{ pc}$, with a robust 2σ upper limit of $r_{\text{e,maj}} < 51 \text{ pc}$.
- We do not detect any significant X-ray emission in either soft or hard X-ray bands (rest $\sim 1\text{--}10 \text{ keV}$ and $10\text{--}30 \text{ keV}$).

We perform spectrophotometric fitting to rigorously test current state-of-the-art stellar population, AGN, and composite models.

- We demonstrate that SED models in which a galaxy (i.e. evolved stellar population) dominates the rest-optical continuum systematically overpredicts the flux blueward, and underpredicts the flux redward of the Balmer break. This remains the case even when including an AGN component in the model, and when allowing for extremely steep dust laws with unusually high optical depths.
- The ill-fitting star-dominated SED models favour a massive ($\log(M_*/M_\odot) \sim 10^{10.4-10.6}$), reddened post-starburst solution. We show that the implied extremely high stellar mass density ($\rho \sim 10^5 M_\odot \text{pc}^{-3}$) would lead to frequent stellar collisions between giant and main-sequence stars with a collision rate of $\sim 5 \text{ yr}^{-1}$, which may produce detectable bursts of X-ray emission.
- By exploring stellar spectra, we demonstrate that main-sequence A stars alone cannot produce the Balmer break and shape of the rest-optical SED, even after attenuation with steep dust laws. Altering the IMF slope to boost the relative number of A stars is thus not sufficient to explain the SED. Only an extremely top-heavy IMF with a significant fraction of supergiants – in combination with a strong and steep dust attenuation curve, and a highly contrived star formation history – can simultaneously produce a strong Balmer break and red rest-optical continuum. But, we estimate that such an IMF would also lower the stellar mass density by >1 dex, and reduce the stellar velocity dispersion by a factor of >3 .

We therefore conclude that the rest-optical and near-IR continuum of *The Cliff* cannot originate from a massive, evolved stellar population with an extremely high stellar density. Crucially, this is the first LRD for which such a strong conclusion can be drawn solely by analysing the SED.

Instead, we argue that the most plausible model is that of a luminous ionising source reddened by dense, absorbing gas in its close vicinity. Currently the only model capable of producing both the strength and shape of the observed Balmer break is that of a black hole star, in which an AGN accretion disc is embedded in turbulent dense gas. Such a model would also naturally explain the observed absorption in the $H\alpha$ line. The BH* models are still in an early phase of development. We show that recently proposed dust-reddened BH* models (with $n_H = 10^{10} \text{ cm}^{-3}$ and $N_H = 10^{24} \text{ cm}^{-2}$) can provide a good fit to the rest-optical, but severely overpredict the rest near- and mid-IR. We therefore propose that the AGN spectrum may be intrinsically red, i.e. intrinsically luminous in the rest-optical in addition to strong reddening from very dense gas ($n_H \sim 10^{11} \text{ cm}^{-3}$; $N_H \sim 10^{26} \text{ cm}^{-2}$), and show that such a BH* model currently provides the best match to the rest-optical to mid-IR of *The Cliff*. Possibly, an intrinsically red spectrum could be explained by less conventional accretion disc models, for instance by invoking super-Eddington accretion (which in turn would imply that our black hole mass estimate is highly uncertain). Alternatively, (super)massive stars formed within the dense gaseous envelope may also contribute to the red continuum emission, but models for such a nuclear star cluster scenario are yet to be explored.

The Cliff presents the strongest direct evidence to date that the Balmer break and rest optical to near-IR SED in LRDs can be dominated by emission from an AGN, rather than evolved stellar populations, although many open questions regarding the black hole and host galaxy properties remain. Because of its comparatively modest redshift, the high-quality spectrophotometric coverage of JWST extends over a wide rest-frame wavelength range. These stringent constraints make *The Cliff* the ideal benchmark for future AGN and BH* models.

Acknowledgements. We thank the PRIMER team for making their imaging data publicly available immediately. We thank Jaime Villaseñor and Friedrich Röpke for helpful discussions. This work is based on observations made with the NASA/ESA/CSA James Webb Space Telescope. The data were obtained from the Mikulski Archive for Space Telescopes at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST. These observations are associated with programs #1837 and #4233. Support for program #4233 was provided by NASA through a grant from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127. REH acknowledges support by the German Aerospace Center (DLR) and the Federal Ministry for Economic Affairs and Energy (BMWi) through program 50OR2403 ‘RUBIES’. This research was supported by the International Space Science Institute (ISSI) in Bern, through ISSI International Team project #562. The Cosmic Dawn Center is funded by the Danish National Research Foundation (DNRF) under grant #140. This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number MB22.00072, as well as from the Swiss National Science Foundation (SNSF) through project grant 200020_207349. Support for this work for RPN was provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51515.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. TBM was supported by a CIERA fellowship.

References

- Abramowicz, M. A., Calvani, M., & Nobili, L. 1980, *ApJ*, **242**, 772
- Akins, H. B., Casey, C. M., Lambrides, E., et al. 2024, *ApJ*, submitted [arXiv:2406.10341]
- Akins, H. B., Casey, C. M., Berg, D. A., et al. 2025a, *ApJ*, **980**, L29
- Akins, H. B., Casey, C. M., Chisholm, J., et al. 2025b, *ApJL*, submitted [arXiv:2503.00998]
- Ananna, T. T., Bogdán, Á., Kovács, O. E., Natarajan, P., & Hickox, R. C. 2024, *ApJ*, **969**, L18
- Baggen, J. F. W., van Dokkum, P., Labbé, I., et al. 2023, *ApJ*, **955**, L12
- Baggen, J. F. W., van Dokkum, P., Brammer, G., et al. 2024, *ApJ*, **977**, L13
- Balogh, M. L., Morris, S. L., Yee, H. K. C., Carlberg, R. G., & Ellingson, E. 1999, *ApJ*, **527**, 54
- Barro, G., Pérez-González, P. G., Kocevski, D. D., et al. 2024a, *ApJ*, **963**, 128
- Barro, G., Pérez-González, P. G., Kocevski, D. D., et al. 2024b, arXiv e-prints [arXiv:2412.01887]
- Barrufet, L., Oesch, P. A., Marques-Chaves, R., et al. 2025, *MNRAS*, **537**, 3453
- Bartko, H., Martins, F., Trippe, S., et al. 2010, *ApJ*, **708**, 834
- Begelman, M. C., Volonteri, M., & Rees, M. J. 2006, *MNRAS*, **370**, 289
- Begelman, M. C., Rossi, E. M., & Armitage, P. J. 2008, *MNRAS*, **387**, 1649
- Bellovary, J. 2025, *ApJ*, **984**, L55
- Bezanson, R., van Dokkum, P. G., Tal, T., et al. 2009, *ApJ*, **697**, 1290
- Binney, J., & Tremaine, S. 2008, *Galactic Dynamics*, 2nd edn. (Princeton, NJ: USA: Princeton University Press)
- Bogdán, Á., Goulding, A. D., Natarajan, P., et al. 2024, *Nat. Astron.*, **8**, 126
- Boylan-Kolchin, M. 2023, *Nat. Astron.*, **7**, 731
- Brammer, G. 2023, <https://doi.org/10.5281/zenodo.1146904>
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
- Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, *ApJ*, **533**, 682
- Carnall, A. C., McLeod, D. J., McLure, R. J., et al. 2023, *MNRAS*, **520**, 3974
- Carnall, A. C., Cullen, F., McLure, R. J., et al. 2024, *MNRAS*, **534**, 325
- Chabrier, G. 2003, *PASP*, **115**, 763
- Charlot, S., & Fall, S. M. 2000, *ApJ*, **539**, 718
- Chevallard, J., Charlot, S., Wandelt, B., & Wild, V. 2013, *MNRAS*, **432**, 2061
- Choi, J., Dotter, A., Conroy, C., et al. 2016, *ApJ*, **823**, 102
- Conroy, C., & Gunn, J. E. 2010, *ApJ*, **712**, 833
- Conroy, C., Gunn, J. E., & White, M. 2009, *ApJ*, **699**, 486
- Coughlin, E. R., & Begelman, M. C. 2024, *ApJ*, **970**, 158
- Cueto, E. R., Hutter, A., Dayal, P., et al. 2024, *A&A*, **686**, A138
- de Graaff, A., Rix, H.-W., Carniani, S., et al. 2024, *A&A*, **684**, A87
- de Graaff, A., Brammer, G., Weibel, A., et al. 2025a, *A&A*, **697**, A189
- de Graaff, A., Setton, D. J., Brammer, G., et al. 2025b, *Nat. Astron.*, **9**, 280
- D’Eugenio, F., Maiolino, R., Perna, M., et al. 2025, *MNRAS*, submitted [arXiv:2503.11752]
- Donnan, C. T., McLure, R. J., Dunlop, J. S., et al. 2024, *MNRAS*, **533**, 3222
- Dotter, A. 2016, *ApJS*, **222**, 8
- Draine, B. T., & Li, A. 2007, *ApJ*, **657**, 810
- Drout, M. R., Massey, P., Meynet, G., Tokarz, S., & Caldwell, N. 2009, *ApJ*, **703**, 441
- Ferland, G. J., Chatzikos, M., Guzmán, F., et al. 2017, *Rev. Mex. Astron. Astrofis.*, **53**, 385

- Fouesneau, M. 2025, <https://doi.org/10.5281/zenodo.14712174>
- Furtak, L. J., Labbé, I., Zitrin, A., et al. 2024, *Nature*, **628**, 57
- Gezari, S. 2021, *ARA&A*, **59**, 21
- Glazebrook, K., Nanayakkara, T., Schreiber, C., et al. 2024, *Nature*, **628**, 277
- Goodman, J. 2003, *MNRAS*, **339**, 937
- Gordon, K. D. 2024, *J. Open Source Softw.*, **9**, 7023
- Gordon, K. D., Clayton, G. C., Misselt, K. A., Landolt, A. U., & Wolff, M. J. 2003, *ApJ*, **594**, 279
- Gordon, K. D., Cartledge, S., & Clayton, G. C. 2009, *ApJ*, **705**, 1320
- Goulding, A. D., Greene, J. E., Setton, D. J., et al. 2023, *ApJ*, **955**, L24
- Greene, J. E., & Ho, L. C. 2005, *ApJ*, **630**, 122
- Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, *ApJ*, **964**, 39
- Grudić, M. Y., Hopkins, P. F., Quataert, E., & Murray, N. 2019, *MNRAS*, **483**, 5548
- Guia, C. A., Pacucci, F., & Kocevski, D. D. 2024, *Res. Notes Am. Astron. Soc.*, **8**, 207
- Habouzit, M. 2025, *MNRAS*, **537**, 2323
- Harikane, Y., Zhang, Y., Nakajima, K., et al. 2023a, *ApJ*, **959**, 39
- Harikane, Y., Ouchi, M., Oguri, M., et al. 2023b, *ApJS*, **265**, 5
- Heintz, K. E., Brammer, G. B., Watson, D., et al. 2025, *A&A*, **693**, A60
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, **208**, 19
- Hoffman, M. D., Gelman, A., et al. 2014, *J. Mach. Learn. Res.*, **15**, 1593
- Hopkins, P. F., Murray, N., Quataert, E., & Thompson, T. A. 2010, *MNRAS*, **401**, L19
- Hu, B. X., & Loeb, A. 2024, *A&A*, **690**, A130
- Hviding, R. E., de Graaff, A., Miller, T. B., et al. 2025, *A&A*, submitted [arXiv:2506.05459]
- Inayoshi, K., & Maiolino, R. 2025, *ApJ*, **980**, L27
- Inayoshi, K., Kimura, S., & Noda, H. 2024, *PASJ*, accepted [arXiv:2412.03653]
- Ji, X., Maiolino, R., Übler, H., et al. 2025, *MNRAS*, submitted [arXiv:2501.13082]
- Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2014, *ApJ*, **796**, 106
- Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2019, *ApJ*, **880**, 67
- Johnson, B., Foreman-Mackey, D., Sick, J., et al. 2021, <https://doi.org/10.5281/zenodo.4577191>
- Juodžbalis, I., Ji, X., Maiolino, R., et al. 2024, *MNRAS*, **535**, 853
- Killi, M., Watson, D., Brammer, G., et al. 2024, *A&A*, **691**, A52
- Kocevski, D. D., Hasinger, G., Brightman, M., et al. 2018, *ApJS*, **236**, 48
- Kocevski, D. D., Finkelstein, S. L., Barro, G., et al. 2025, *ApJ*, **986**, 126
- Kokorev, V., Fujimoto, S., Labbe, I., et al. 2023, *ApJ*, **957**, L7
- Kokorev, V., Caputi, K. I., Greene, J. E., et al. 2024, *ApJ*, **968**, 38
- Labbé, I., van Dokkum, P., Nelson, E., et al. 2023, *Nature*, **616**, 266
- Labbe, I., Greene, J. E., Matthee, J., et al. 2024, *ApJ*, submitted [arXiv:2412.04557]
- Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, *ApJ*, **978**, 92
- Lambrides, E., Garofali, K., Larson, R., et al. 2024, arXiv e-prints [arXiv:2409.13047]
- Laor, A. 2006, *ApJ*, **643**, 112
- Leja, J., Johnson, B. D., Conroy, C., van Dokkum, P. G., & Byler, N. 2017, *ApJ*, **837**, 170
- Leja, J., Johnson, B. D., Conroy, C., & van Dokkum, P. 2018, *ApJ*, **854**, 62
- Leung, G. C. K., Finkelstein, S. L., Pérez-González, P. G., et al. 2024, *ApJ*, submitted [arXiv:2411.12005]
- Li, J., Silverman, J. D., Shen, Y., et al. 2025, *ApJ*, **981**, 19
- Libralato, M., Argyriou, I., Dicken, D., et al. 2024, *PASP*, **136**, 034502
- Lu, J. R., Do, T., Ghez, A. M., et al. 2013, *ApJ*, **764**, 155
- Lupi, A., Trinca, A., Volonteri, M., Dotti, M., & Mazzucchelli, C. 2024, *A&A*, **689**, A128
- Ma, Y., Greene, J. E., Setton, D. J., et al. 2025, *ApJ*, **981**, 191
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, *A&A*, **691**, A145
- Martins, F., Schaerer, D., Haemmerlé, L., & Charbonnel, C. 2020, *A&A*, **633**, A9
- Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, *ApJ*, **963**, 129
- Matthee, J., Naidu, R. P., Kotiwale, G., et al. 2025, *ApJ*, **988**, 246
- Mauerhofer, V., Dayal, P., Haehnelt, M. G., et al. 2025, *A&A*, **696**, A157
- Mérida, R. M., Gaspar, G., Sawicki, M., et al. 2025, *A&A*, **698**, A317
- Naidu, R. P., Matthee, J., Katz, H., et al. 2025, arXiv e-prints [arXiv:2503.16596]
- Nanayakkara, T., Glazebrook, K., Schreiber, C., et al. 2025, *ApJ*, **981**, 78
- Narayanan, D., Conroy, C., Davé, R., Johnson, B. D., & Popping, G. 2018, *ApJ*, **869**, 70
- Neškova, M., Sirocky, M. M., Nikutta, R., Ivezić, Ž., & Elitzur, M. 2008, *ApJ*, **685**, 160
- Neumayer, N., Seth, A., & Böker, T. 2020, *A&ARv*, **28**, 4
- Noll, S., Burgarella, D., Giovannoli, E., et al. 2009, *A&A*, **507**, 1793
- Pacucci, F., Nguyen, B., Carniani, S., Maiolino, R., & Fan, X. 2023, *ApJ*, **957**, L3
- Pasha, I., & Miller, T. B. 2023, *J. Open Source Softw.*, **8**, 5703
- Pérez-González, P. G., Barro, G., Rieke, G. H., et al. 2024, *ApJ*, **968**, 4
- Phan, D., Pradhan, N., & Jankowiak, M. 2019, arXiv e-prints [arXiv:1912.11554]
- Pizzati, E., Hennawi, J. F., Schaye, J., et al. 2025, *MNRAS*, **539**, 2910
- Portegies Zwart, S. F., & McMillan, S. L. W. 2002, *ApJ*, **576**, 899
- Rigby, J., Perrin, M., McElwain, M., et al. 2023, *PASP*, **135**, 048001
- Roberts-Borsani, G., Treu, T., Shapley, A., et al. 2024, *ApJ*, **976**, 193
- Rusakov, V., Watson, D., Nikopoulos, G. P., et al. 2025, *Nature*, submitted [arXiv:2503.16595]
- Ryu, T., Amaro Seoane, P., Taylor, A. M., & Ohlmann, S. T. 2024, *MNRAS*, **528**, 6193
- Salim, S., & Narayanan, D. 2020, *ARA&A*, **58**, 529
- Salim, S., Boquien, M., & Lee, J. C. 2018, *ApJ*, **859**, 11
- Sánchez-Blázquez, P., Peletier, R. F., Jiménez-Vicente, J., et al. 2006, *MNRAS*, **371**, 703
- Setton, D. J., Greene, J. E., de Graaff, A., et al. 2024, *ApJ*, submitted [arXiv:2411.03424]
- Setton, D. J., Greene, J. E., Spilker, J. S., et al. 2025, *ApJL*, submitted [arXiv:2503.02059]
- Shen, X., Hopkins, P. F., Faucher-Giguère, C.-A., et al. 2020, *MNRAS*, **495**, 3252
- Sirko, E., & Goodman, J. 2003, *MNRAS*, **341**, 501
- Speagle, J. S. 2020, *MNRAS*, **493**, 3132
- Steinhardt, C. L., Kokorev, V., Rusakov, V., Garcia, E., & Sneppen, A. 2023, *ApJ*, **951**, L40
- Taylor, A. J., Finkelstein, S. L., Kocevski, D. D., et al. 2025, *ApJ*, **986**, 165
- Temple, M. J., Hewett, P. C., & Banerji, M. 2021, *MNRAS*, **508**, 737
- Valentino, F., Brammer, G., Gould, K. M. L., et al. 2023, *ApJ*, **947**, 20
- van Dokkum, P., & Conroy, C. 2024, *ApJ*, **973**, L32
- Volonteri, M., & Begelman, M. C. 2010, *MNRAS*, **409**, 1022
- Wang, B., Leja, J., de Graaff, A., et al. 2024a, *ApJ*, **969**, L13
- Wang, B., Leja, J., Labbé, I., et al. 2024b, *ApJS*, **270**, 12
- Wang, B., de Graaff, A., Davies, R. L., et al. 2025, *ApJ*, **984**, 121
- Watanabe, S., & Oppen, M. 2010, *J. Mach. Learn. Res.*, **11**, 3571
- Weibel, A., Oesch, P. A., Barrufet, L., et al. 2024, *MNRAS*, **533**, 1808
- Weibel, A., de Graaff, A., Setton, D. J., et al. 2025, *ApJ*, **983**, 11
- Williams, C. C., Alberts, S., Ji, Z., et al. 2024, *ApJ*, **968**, 34
- Woodrum, C., Rieke, M., Ji, Z., et al. 2024, *Proc. Natl. Acad. Sci.*, **121**, e2317375121
- Xiao, M., Oesch, P. A., Bing, L., et al. 2025, *A&A*, **700**, A231
- Yue, M., Eilers, A.-C., Ananna, T. T., et al. 2024, *ApJ*, **974**, L26

Appendix A: JWST/NIRCam and MIRI photometry

We show the JWST images, MAP models from `pysersic` (Section 2.1), and image residuals (computed as data–model) of all NIRCam and MIRI images. Specifically, Figure A.1 shows the MIRI image cutouts and results from the morphological fitting. Figure A.2 presents the F200W NIRCam image, model and residual for the higher pixel scale of $0.02'' \text{ pix}^{-1}$ used in Section 4. The NIRCam images and fitting results at pixel scale of $0.04'' \text{ pix}^{-1}$ are shown in Figure A.3 for all 8 NIRCam filters. We find that the Sérsic profile fitting performs well across all wavelengths, with only minor residual features remaining.

We use these models to determine the total integrated flux of *The Cliff*, presented in Table A.1. Moreover, we quantify to what extent the flux at the position of the NIRSpec microshutter aperture originates from *The Cliff* versus the foreground source. To do so, we first measure the shutter aperture flux from the individual NIRCam images. We account for uncertainty in the pointing accuracy of JWST by applying random spatial shifts, using 500 draws from a Gaussian distribution with a dispersion of $0.025''$ (Rigby et al. 2023). Next, we perform the same aperture measurements (including the same spatial shifts) for the model image of *The Cliff* alone; to at the same time account for uncertainties in the morphological fitting, we generate 500 model images drawn from the posteriors of the `pysersic` fits. The ratio of these aperture fluxes is shown for the 8 NIRCam filters in Figure A.4. *The Cliff* clearly dominates the emission across all wavelengths ($> 80\%$), and accounts for 100% of the flux redward of the Balmer limit.

We therefore conclude that there is no contamination from the foreground source within the NIRSpec aperture at $> 2 \mu\text{m}$. At $< 2 \mu\text{m}$ there may be slight contamination (at the $< 20\%$ level), but this does not significantly affect any of the modelling or conclusions – which rely on the flux at $> 2 \mu\text{m}$ – in this work. The Balmer break strength could be even higher than measured

from the spectrum, by at most 25%, but this is subdominant to the large measurement uncertainty shown in Figure 2.

Table A.1. NIRCam and MIRI photometry of *The Cliff* obtained with `pysersic` (Section 2.1).

filter	f_ν (μJy)
NIRCam/F090W	$0.0148^{+0.0024}_{-0.0019}$
NIRCam/F115W	$0.028^{+0.004}_{-0.004}$
NIRCam/F150W	$0.0269^{+0.0022}_{-0.0017}$
NIRCam/F200W	$0.322^{+0.029}_{-0.028}$
NIRCam/F277W	$1.090^{+0.004}_{-0.004}$
NIRCam/F356W	$1.334^{+0.005}_{-0.006}$
NIRCam/F410M	$1.597^{+0.011}_{-0.013}$
NIRCam/F444W	$1.725^{+0.012}_{-0.013}$
MIRI/F770W	$1.54^{+0.03}_{-0.03}$
MIRI/F1800W	$1.96^{+0.5}_{-0.5}$




Fig. A.1. JWST/MIRI images of *The Cliff* and its nearby lower-redshift neighbour for all available filters. We show the MAP models from the point source + Sérsic profile fitting for each filter as well as the residuals.




Fig. A.2. JWST/NIRCam F200W image at $0.02'' \text{ pix}^{-1}$ of *The Cliff*, its nearby lower-redshift neighbour, and the nearby star. We show the MAP model as well as the residuals.

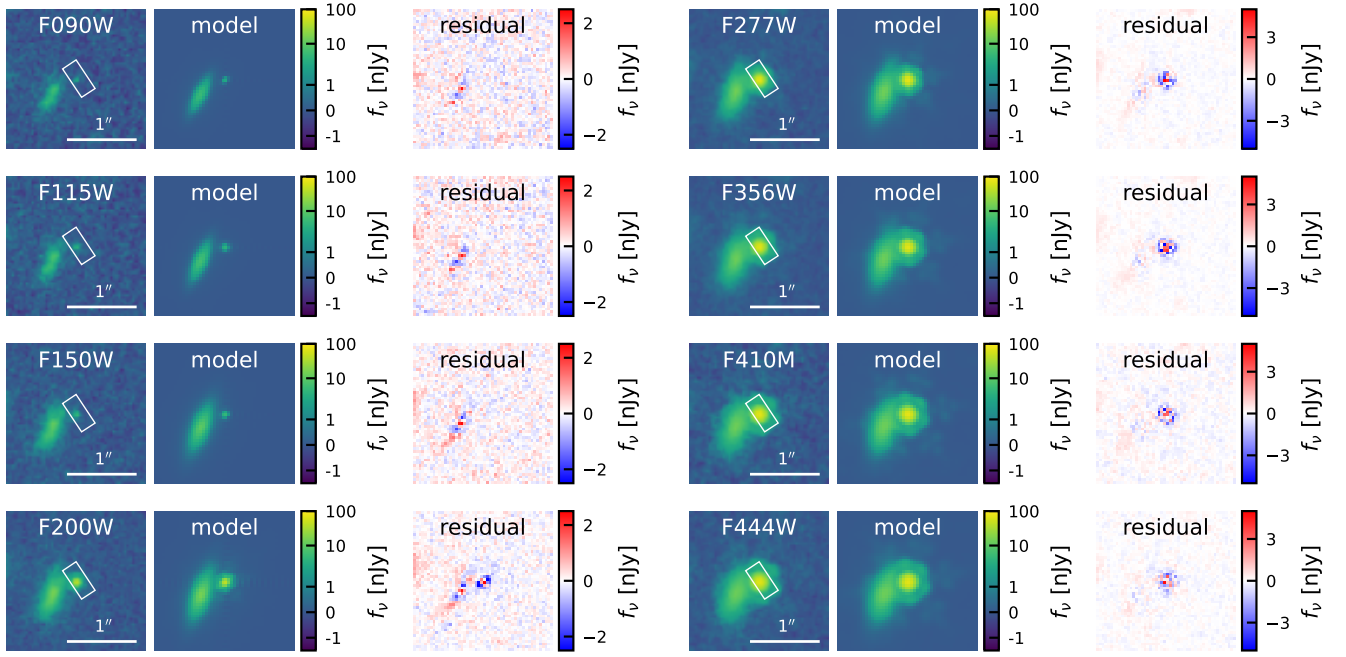


Fig. A.3. JWST/NIRCam images ($0.04'' \text{ pix}^{-1}$) of *The Cliff* and its nearby lower-redshift neighbour for all available filters. We show the MAP models from the Sérsic profile fitting for each filter as well as the residuals. The central NIRSpec shutter is plotted for reference.




Fig. A.4. Ratio of the flux measured from the Sérsic model of *The Cliff* to the observed NIRCam image within the rectangular aperture of the NIRSpec microshutter. Error bars reflect both the uncertainty in the Sérsic fitting and the exact pointing location of JWST. Emission from *The Cliff* dominates ($> 80\%$) the total observed flux at all wavelengths, and accounts for $\sim 100\%$ of the total emission redward of the Balmer break ($> 2\mu\text{m}$).